DataGeneral
Data General Corporation, Westboro, Massachusetts 01580

Customer Documentation

Managing the DG/UX™ System

093-701088-04

AVIION®
PRODUCT LINE

Managing the DG/UX™ System

093–701088–04

For the latest enhancements, cautions, documentation changes, and other information on this product, please see the Release Notice (085–series) and/or Update Notice (078–series) supplied with the software.

Printed in the United States of America
Rev. 04, *January*, 1994

Ordering No. 093–701088–04

Notice

Managing the DG/UX™ System
093-701088-04

Revision History:		Effective with:
Original Release	– June, 1991	
Revision 1	– January, 1992	
Revision 2	– August, 1992	
Revision 3	– March, 1993	DG/UX System 5.4 Release 2.01
Revision 4	– January, 1994	DG/UX System 5.4 Release 3.00

A vertical bar in the page margin shows major technical change from the previous revision. The exceptions are Chapters 7 and 14, which are substantially revised, and Chapters 8 and 16 and Appendix D, which are new.

Preface

This manual introduces DG/UX™ system management and provides complete coverage of the various DG/UX system-related tasks that the typical system administrator must address. The manual does not cover installation. To install your DG/UX operating system software, see *Installing the DG/UX™ System*. For a streamlined treatment of the tasks you perform to make a freshly-installed system usable, see *Customizing the DG/UX™ System*.

The DG/UX system provides the **sysadm**(1M) utility to help you with installation and management tasks. The **sysadm** utility is a menu-based interface to a number of programs intended to simplify the responsibilities of a system administrator. Using **sysadm** and its attendant programs, you can do things such as load and set up software packages, add user profiles, manage the network and printers, and build customized DG/UX kernels. This manual tells how you can use **sysadm** to help with these and other tasks.

Manual Outline

Many of the chapters in this manual are divided into two distinct sections. First is a section containing all the basic information needed to administer some part of the DG/UX system. Next is an optional "Expert Information" section containing more advanced information. You can use these sections to enhance your knowledge of the DG/UX system.

The manual contains the following chapters and appendixes:

Chapter 1 **Introduction to System Management**
Outlines some of the more common duties of an administrator. Includes a description of the **sysadm** utility and some common system configurations.

Chapter 2 **Overview of Routine Management Tasks**
Covers tasks that you should perform every time you boot the system and then periodically while the system is running.

Chapter 3 **Operating the DG/UX System**
Explains startup and shutdown, recovering from system trouble (such as crashes and power outages). Explains run command scripts and run levels.

Chapter 4 **Managing the System Configuration**
Covers setting system parameters and time and date, configuring the system, and rebuilding the kernel on a routine basis. There is a new section that explains DG/UX system user licenses, how to get information about a user license, and how to upgrade the user count for a user license.

Chapter 5 **Managing Software**
Covers creating releases, loading and setting up software packages.

Chapter 6 **Managing OS and X Terminal Clients**
Covers setting defaults, adding and deleting OS and X terminal clients.

Chapter 7 **Managing Disks**
Using **sysadm**, introduces virtual disk management, a new technology that enables online disk operations, which do not require that file systems be unmounted first. Describes how to format physical disks, create virtual disks and file systems, check file systems, change file system size, display disk information, update the operating system, and do various other disk-related tasks. The chapter also covers special disk and file system-related features such as mirrored virtual disks, cached virtual disks, data striping, fast recovery file systems, and write verification.

Chapter 8 **Setting Up and Managing Features of High–Availability**
Explains three categories of configurations to increase high availability: disk failover, IP takeover and NFS failover, and multi–path LAN I/O.

Chapter 9 **Managing File Systems**
Shows how to mount, unmount, add, delete, shrink, and expand file systems, manage swap areas, list information about file systems, search for possible security problems, make backup tapes, and check the consistency of file systems.

Chapter 10 **Managing Terminals and Ports**
Describes how to setup and manage terminals and ports.

Chapter 11 **Managing Controllers**
Describes how to manage LAN, Sync, and VDA controllers.

Chapter 12 **Managing Printers**
Contains operations for managing printing devices, classes, filters, forms, requests, and the scheduler.

Chapter 13 **Managing Login Accounts**
Covers adding login accounts, creating aliases and groups. Explains how these are used and managed differently in stand-alone and OS server-client environments.

Chapter 14 **Managing Accounting**
Monitoring the use of system resources using the accounting facility; printing various kinds of summary reports.

Chapter 15 **Using CD-ROM, Magneto-optical, Diskette, and Tape Drives**
Describes special requirements for SCSI devices.

Chapter 16 **Logging System and Network Errors**
Explains how to maintain a record of system and network errors.

Appendix A **Fsck Error Conditions**
Covers file system checker (**fsck**(1M)) error conditions.

Appendix B **SAF Reference Cards**
Contains reference tables summarizing commands used with SAF, which is discussed in Chapter 10.

Appendix C **DG/UX Directories and Files**
Selective summary of directories and files of interest to the system administrator.

Appendix D **SCSI Disk and Tape Drive Names**
Outlines SCSI disk and tape drive and controller names, including disk-array storage system drive and controller names.

Readers, Please Note

Data General manuals use certain symbols and styles of type to indicate different meanings. The Data General symbol and typeface conventions used in this manual are defined in the following list. You should familiarize yourself with these conventions before reading the manual.

This manual also presumes the following meanings for the terms "command line," "format line," and "syntax line." A command line is an example of a command string that you should type verbatim; it is preceded by a system prompt and is followed by a delimiter such as the curved arrow symbol for the New Line key. A format line shows how to structure a command; it shows the variables that must be supplied and the available options. A syntax line is a fragment of program code that shows how to use a particular routine; some syntax lines contain variables.

Convention	Meaning
boldface	In command lines and format lines: Indicates text (including punctuation) that you type verbatim from your keyboard. All DG/UX commands, pathnames, and names of files, directories, and manual pages also use this typeface.
`constant width/ monospace`	Represents a system response on your screen. Syntax lines also use this font.
italic	In format lines: Represents variables for which you supply values; for example, the names of your directories and files, your username and password, and possible arguments to commands. In text: Indicates a term that is defined in the manual's glossary.
[ ]	In format lines: These brackets surround an optional argument. Don't type the brackets; they only set off what is optional. The brackets are in regular type and should not be confused with the boldface brackets shown below.
[ ]	In format lines: Indicates literal brackets that you should type. These brackets are in boldface type and should not be confused with the regular type brackets shown above.

...	In format lines and syntax lines: Means you can repeat the preceding argument as many times as desired.
$ and %	In command lines and other examples: Represent the system command prompt symbols used for the Bourne and C shells, respectively. Note that your system might use different symbols for the command prompts.
⤶	In command lines and other examples: Represents the Enter key, which is the name of the key used to generate a new line. (Note that on some keyboards this key might be called New Line or Return instead of Enter.) Throughout this manual, a space precedes the Enter symbol to improve readability; you can ignore it.
< >	In command lines and other examples: Angle brackets distinguish a command sequence or a keystroke (such as **<Ctrl–D>**, **<Esc>**, and **<3dw>**) from surrounding text. Note that these angle brackets are in regular type and that you do not type them; there are, however, boldface versions of these symbols (described below) that you do type.
<, >, >>	In text,command lines, and other examples: These boldface symbols are redirection operators, used for redirecting input and output. When they appear in boldface type, they are literal characters that you should type.
▯	In command lines and other examples: Represents the cursor, which indicates your current typing position on the screen.

Contacting Data General

Data General wants to assist you in any way it can to help you use its products. Please feel free to contact the company as outlined below.

Manuals

If you require additional manuals, please use the enclosed TIPS order form (United States only) or contact your local Data General sales representative.

Telephone Assistance

If you are unable to solve a problem using any manual you received with your system, free telephone assistance is available with your hardware warranty and with most Data General software service options. If you are within the United States or Canada, contact the Data General Customer Support Center (CSC) by calling 1–800–DG–HELPS. Lines are open from 8:00 a.m. to 5:00 p.m., your time, Monday through Friday. The center will put you in touch with a member of Data General's telephone assistance staff who can answer your questions.

For telephone assistance outside the United States or Canada, ask your Data General sales representative for the appropriate telephone number.

Joining Our Users Group

Please consider joining the largest independent organization of Data General users, the North American Data General Users Group (NADGUG). In addition to making valuable contacts, members receive FOCUS monthly magazine, a conference discount, access to the Software Library and Electronic Bulletin Board, an annual Member Directory, Regional and Special Interest Groups, and much more. For more information about membership in the North American Data General Users Group, call 1–800–253–3902 or 1–508–443–3330.

End of Preface

Contents

Chapter 1 – Introduction to System Management

Chapter 2 – Overview of Routine Management Tasks

Chapter 3 – Operating the DG/UX System

Chapter 4 – Managing the System Configuration

Chapter 5 – Managing Software

Chapter 6 – Managing OS and X Terminal Clients

Chapter 7 – Managing Disks

Chapter 8 – Setting Up and Managing Features of High Availability

Chapter 9 – Managing File Systems

Chapter 10 – Managing Terminals and Ports

Chapter 11 – Controller Management

Chapter 12 – Printer Management

Chapter 13 – Login Account Management

Chapter 14 – Accounting

Chapter 15 – Using CD-ROM, Magneto-Optical, Diskette, and Tape Drives

Chapter 16 – Logging System Errors

Appendix A – fsck Error Conditions

Appendix B– SAF Reference Cards

Appendix C – DG/UX Directories and Files

Appendix D – SCSI Disk and Tape Drive Names

Tables

Table

Figures

Figure

Chapter 1
Introduction to System Management

This chapter introduces you to the role of system administrator, identifying the basic tasks you will perform on a regular basis and the tools you use to effectively administer the system. Also, it introduces the typical types of configurations you may be administering.

Duties of the System Administrator

Your goal as system administrator is to keep your system (hardware and software) running smoothly so that users can get their work done. You will perform some tasks on a daily basis such as monitoring and cleaning up the file system. Other duties you will perform periodically; often when a particular situation arises (such as adding new user accounts, configuring new peripherals, or shutting down the system to perform hardware maintenance). From the extensive system management tasks discussed in this manual, you can pick those that apply to your environment; thus, developing your own system management style.

Read a summary of the manual's outline in the Preface and scan the table of contents of this manual to get an idea of the nature of system management tasks. A condensed list follows of the more common tasks:

- Adding user accounts and responding to user questions and problems.
- Maintaining hardware (connecting new computers, printers, terminals, disk drives, modems, installing new boards, making network connections, and ensuring their configuration).
- Installing and configuring new operating system and additional software package releases and updates.
- Monitoring the system's health through managing disk space, maintaining file systems, backing up data, and checking log file messages.
- Troubleshooting hardware and software problems and performing diagnostic procedures such as taking the system down, performing system dumps and system backups, and rebooting the system.
- Keeping the network services operational, including electronic mail.
- Ensuring system security (assigning file and directory permissions, and user passwords).

Logging in as Superuser

As system administrator, you will execute operations and commands which are restricted to the superuser whose default logins are **sysadm** and **root**. You should use the **sysadm** rather than the **root** profile when performing administrative tasks, so that commands executed can be performed and output generated in the **/admin** directory rather than the / directory, which is the home directory of the **root** profile.

To become superuser, either log in as the superuser using the **sysadm** login, or if already logged in, use the **su command**, specifying **- sysadm** as the desired user profile.

Example:

```
% su - sysadm ↵
```

You are then prompted for the password to the superuser profile. If you have not assigned passwords to the **sysadm** and **root** profiles, you should do so immediately; otherwise, any user may log in as the superuser.

Examples:

```
# passwd sysadm ↵
```

or

```
# passwd root ↵
```

The System Administration (sysadm) Utility

You use the **sysadm** utility to manage your system. To begin DG/UX system administration, become superuser and start **sysadm** as follows:

```
# sysadm ↵
```

If you are working from a workstation or X terminal, the **sysadm** main window appears:

Answering "no" to this prompt does not give you the option to select an alternate mount point.

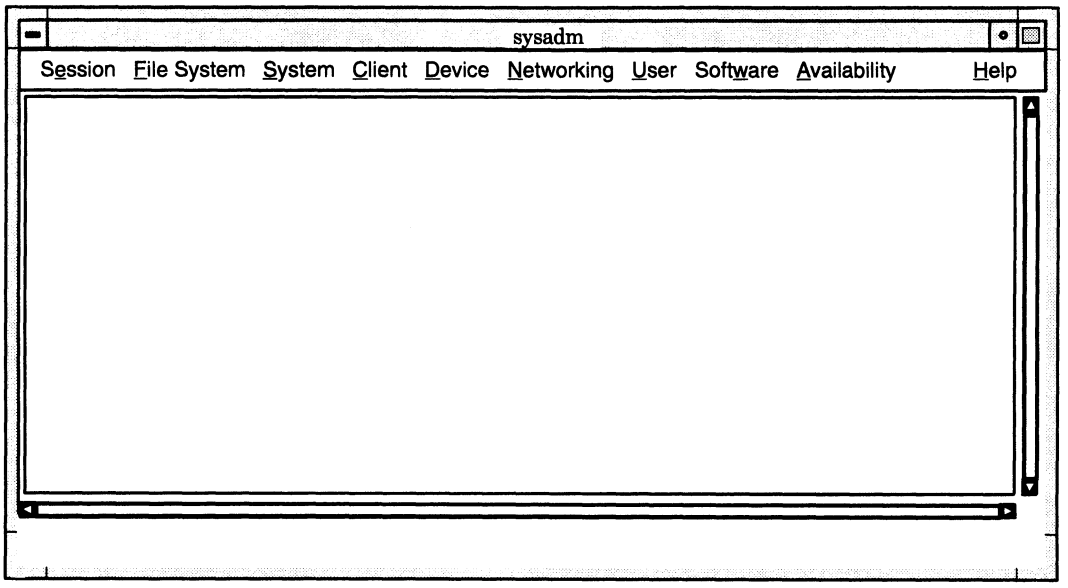

Figure 1–1 The Graphical Version of the sysadm Main Window

The window that actually appears on your system may have more or fewer selections than those shown in Figure 1–1. This is because **sysadm** window and menu displays vary depending on the software that is loaded on the system.

If you are working from an ASCII terminal (or if you enter **asysadm** instead of **sysadm**), an ASCII version of **sysadm** appears as shown in Figure 1–2:

```
                  Main Menu

  1  Session ->          Manage this sysadm session
  2  File System ->      Manage file systems
  3  System ->           Manage DG/UX system databases
  4  Client ->           Manage OS and X terminal clients
  5  Device ->           Manage devices and device queues
  6  Logging ->          Manage system and network logging
  7  Networking ->       Manage the network
  8  User ->             Manage users and groups
  9  Software ->         Manage software packages
 10  Availability ->     Manage high availability features
 11  Help ->             Get help on sysadm and its queries

Enter a number, a name, ? or <number>? for help, <NL> to redisplay
menu,or q to quit:
```

Figure 1–2 The ASCII Version of the sysadm Main Menu

The graphical and ASCII interfaces are functionally equivalent. Illustrations in this manual are taken from the ASCII version of **sysadm**.

Selecting Menu Options

In the graphical interface, you select with the:

Mouse	With the mouse pointer on your choice, click button 1.
Keyboard	From the Main window, type **alt–***letter*, where *letter* is the underlined letter (**alt–u** for User, for example). From a pull–down window, type *letter* (**a** to select User Account on the User pull–down menu, for example).
Arrow keys	From the User pull–down, with a box around User Account, press the right arrow key to see the User Account functions.

In the ASCII interface, you can select menu options by entering the number of a menu item or by typing an abbreviation of the menu item name. For example, from the Main Menu, typing "`se`" and pressing Enter selects Session, entering "`sy`" selects System and, entering "`so`" selects Software.

You can start the ASCII interface in the window you want to use. For example, to start **asysadm**, skipping directly to the `User -> Group` menu, enter:

```
# sysadm -m User:Group ↵
```

The system invokes the appropriate interface—the ASCII version (**asysadm**) or graphical version (**xsysadm**) of the **sysadm** utility—based on your display type. If you are working in an X windowing environment, invocation of **sysadm** will produce an X version of **sysadm**; when working in an ASCII terminal environment, by default, you will get an ASCII version of **sysadm**. You may choose explicitly the desired interface, however, by invoking either **xsysadm** or **asysadm**.

In this manual, to illustrate the selection of the List option on the Group menu, we will write: select `User -> Group -> List`. Depending on the interface you're using and your starting point, you interpret the instructions in one of the following ways:

- click User in the **sysadm** main menu, Group in the User pull-down menu, and List in the Group pull-down menu
- enter `8` from the **asysadm** main menu, `2` from the User menu, and `4` from the Group menu
- start the utility with the command:

```
# asysadm -m User:Group:List ↵
```

When repeating the same **sysadm** operation from the graphical interface (for example, adding or changing several dozen accounts), you can save mouse clicks by using a *tear-off menu* rather than the equivalent pull-down menu. There is a dashed line at the top of the User pull-down menu. This indicates that the pull-down menu is optionally a tear-off window. If you select the dashed line, the pull-down menu becomes a tear-off menu at the new location and remains in place on your screen until you close it.

With both a pull-down menu and a tear-off menu, you return to the **sysadm** main menu each time you complete an operation.

Getting Context–sensitive Help

For descriptions of general topics, select Help. Thus, for a description of menu option selection methods and tear-off windows discussed in the previous section, select `Help -> On Interface`.

For help messages describing particular details, use the **sysadm** context-sensitive help facility. At any **sysadm** prompt, you can either enter the requested information or request an explanation. In the graphical interface, press function key 1 (F1) to request a help message. In the ASCII interface, type a question mark (`?`) and press Enter. (In the ASCII interface, obtain a description of a menu by entering *n*`?`, where *n* is the number of the menu.)

For example, when creating a virtual disk, you see the prompt:

```
Stripe Size (in blocks): [16]
```

You can enter a value or you can first request an explanation of the prompt by pressing F1 if you are using the graphical interface or by typing `?` and pressing Enter if you are using the ASCII interface.

Typical Configurations

In discussions of system administration, you will come across terms that refer to the roles played by AViiON® computers running the DG/UX system. In general, a computer that provides network service (such as the operating system or file service) is a server. Correspondingly, computers that receive such services are considered clients. In the DG/UX environment, these are the most common configurations of AViiON computers:

- OS server

  This is an AViiON computer that provides user services through direct terminal lines and/or a local area network (LAN). The server provides operating system (OS) resources through direct terminal access or to client computers across a LAN. In addition to offering OS resources, a server might also provide file service.

  The system administrator may manage the DG/UX system through either a system console (an asynchronous terminal with a keyboard) or a workstation's system console, which consists of the graphics monitor and keyboard. A workstation is an AViiON computer that provides graphical capabilities.

- OS client

  This is an AViiON computer that may or may not have its own disk; thereby relying on the OS server for some or all of its operating system software and file service. Once the OS server establishes a connection with the OS client (by

transferring a kernel image), the OS client subsequently can be responsible for booting its OS over the network and mounting remote file systems.

If an OS client does have an attached disk, it can accommodate some or all of its own operating system. There are various hybrid OS client configurations available for improving system performance. See *Customizing the DG/UX™ System* for more information.

After OS services have been provided to the OS client, the user of the OS client machine assumes certain system administrative responsibilities such as maintenance of the local disk, mounting remote file systems, booting the system, and taking the system down.

- Stand-alone workstation

A workstation is an AViiON computer that provides graphical capabilities. When it stands alone, it has an attached disk subsystem which provides its own OS resources. However, a network connection is often desired for access to remote file systems and communications services such as electronic mail.

End of Chapter

Chapter 2
Overview of Routine Management Tasks

This chapter briefly describes general concerns as well as the day-to-day tasks of system management. In some cases, the discussion includes hints to make the task easier. Topics include:

- Operating policy
- Communicating with users
- Boot time responsibilities
- Day-to-day duties

Operating Policy

Although the virtual disk management (VDM) technology that was introduced in DG/UX 5.4R3.00 offers many new high availability features, situations may arise that require you to shut down the system with little or no notice to users. (See Chapter 7 for more information about VDM.) Try to provide as much advance notice as possible about events affecting the use of the system. When you must take the system out of service, be sure to tell users and OS client managers when the system will be available again. The following section, "Communicating with Users," discusses the various utilities you can use to keep your users informed and alert them to system developments.

You may want to keep the following guidelines in mind before performing any task that will require the system to leave the multiuser state.

- When possible, schedule system maintenance during periods of low system use.
- See if anyone is logged in before taking any actions that affect users. Always give users enough time (at least 60 seconds) to finish whatever they are doing and log off before taking stand-alone or server systems down.

Maintaining a System Log

We recommend that you maintain a detailed system log, both on paper and on disk, of the status of your system. This log may contain the following information:

- What devices are configured into the current kernel. Configured devices are listed in the system file which is located at **/usr/src/uts/aviion/Build/system.***name*, where *name* is usually the host's name.

- Equipment and system configuration changes (dates and actions)
- A record of any system panics and hangs and activities occurring at the time of failure
- Maintenance records (dates and actions)
- A record of problems and solutions

Whatever format you choose for your log, make sure that the system log and the items noted there follow a logical structure. The way you use your system will dictate the form that the logs takes and the diligence with which you maintain it. A system log book can be a valuable tool when trouble-shooting transient problems or when trying to establish system operating characteristics over a period of time.

For problems you encounter after attempting failure recovery, we request that you fill out a Software Trouble Report (STR) and deliver it to the Data General Customer Support Center. For more information on STRs, see the DG/UX system release notice, **/usr/release/dgux_5.4R3.00.rn**

Maintaining a User Trouble Log

As the system administrator, you can expect users to look to you to help solve any number of problems. You may find it useful to log these problems and their solutions in a User Trouble Log not unlike the System Trouble Log you may also keep. With a well-maintained log, you can not only keep track of the solutions to common problems, you may also be able to detect patterns in the problems you encounter, possibly indicating ways in which you can improve service for your users.

Communicating with Users

This section discusses several of the utilities you can use to keep your users informed of system news and administrative developments.

Message of the Day (motd)

You can put items of broad interest that you want to make available to all users in the **/etc/motd** file. The contents of **/etc/motd** are displayed on the user's terminal as part of the login process. The login process executes global files called **/etc/profile** for **sh** and **ksh** users and **/etc/login.csh** for **csh** users. In these files is commonly contained the command:

```
cat /etc/motd
```

Any text contained in **/etc/motd** is displayed for each user each time the user logs in. For this information to have any impact on users, you must take pains to use it sparingly and to clean out outdated announcements. A typical use for the Message of the Day facility might be to publicize down-time:

```
5/30: The system will be down from 6-11pm Thursday for preventive
maintenance.  Call Bob if you have a problem with this.
```

News

The **news** facility, which is an electronic bulletin board, provides a convenient means of distributing announcements to users. The facility maintains a directory, **/usr/news**, where you can put announcements in text files, the names of which are usually used to provide an indication of the content of the news item. The **news** command displays the items on your terminal.

You can also use the **/etc/profile** or **/etc/login.csh** file to inform users about news items. A typical **/etc/profile** contains the line:

```
news -n
```

The `-n` argument causes the names of files in the **/usr/news** directory to be printed on a user's terminal as the user logs in. Item names are displayed only for current items, that is, items added to the **/usr/news** directory since the user last looked at the news. The idea of currency is implemented like this: when you read a news item an empty file named **.news_time** is written in your login directory. As with any other file, **.news_time** carries a time stamp indicating the date and time the file was created. When you log in, the system compares the time stamp of your **.news_time** file and time stamp of items in **/usr/news**.

Unlike **/etc/motd**, where users have no ability to turn the message off, with **news** users have a choice of several possible actions:

read everything
: If the user enters the command, **news** with no arguments, all news items posted since the last time the user typed in the command are printed on the user's terminal.

select some items
: If the **news** command is entered with the names of one or more items as arguments, only those items selected are printed.

read and delete
: After the **news** command has been entered, the user can stop any item from printing by pressing the DELETE key. Pressing the DELETE key twice in a row stops the program.

ignore everything
: If the user is too busy to read announcements at the moment, the messages can safely be ignored. Items remain in **/usr/news** until removed. The item names will continue to be displayed each time the user logs in.

flush all items
: If the user wants simply to eliminate the display of item names without looking at the items, a couple of techniques will work:

  - Use the **touch**(1) command to change the time-accessed and time-modified of the **.news_time** file, thus causing **news** to consider all existing news announcements as having already been read. The following example shows how to use the **touch** command for this purpose:

```
$ touch .news_time ↵
```

- Invoke **news** to read the articles, but direct the output to **/dev/null**. Like the previous technique, this one causes **news** to update the time stamp of the **.new_time** file:

```
$ news > /dev/null ↵
```

Broadcast to All Users

With the **wall**(1M) command, you can broadcast messages to the screens of all users currently logged on the system. While **wall** is a useful device for getting urgent information out quickly, users tend to find it annoying to have messages print out on their screens right in the middle of whatever else is going on. The effect is not destructive, but is somewhat irritating. It is best to reserve this for those times when you need to ask users to log out of the system so that you may perform an administrative task. For example:

```
% wall ↵
The system is going down at 3:00 for oil change -- Marv ↵
<Ctrl-D>
```

In network environments, there is also the **rwall**(1C) command. Use the **rwall** command to send a message to all users on the specified systems. For example:

```
% rwall sales03 accounts lifers ↵
The new desk blotters are in!  Bobbi.
<Ctrl-D>
```

DG/UX System Mail

The DG/UX system has three electronic mail utilities that users can use to communicate among themselves: **mail**(1), **mailx**(1), and **xdgmail**(1X). If your system is connected to others by networking facilities, you can use these utilities to communicate with persons on other systems.

The **mail** program is the basic utility for sending messages. The **mailx** program uses **mail** to send and receive messages, but adds some useful features for storing messages, adding headers, and many other functions. The **xdgmail** program, which also provides a number of extra features, is for workstation users in an OSF/Motif™ environment.

Users can, by default, send and receive mail when you add them to the system with the Add operation. A simple example of using **mailx** follows. In a hypothetical world, Poulet wants to send a mail message to Moe. He types:

```
$ mailx moe ↵
Subject: rubber chicken ↵
Please return my rubber chicken when the puppet show is over. ↵
<Ctrl-D>
```

Poulet enters a subject line when prompted, presses Enter, and then types the message. When finished, he presses **<Ctrl-D>** on the next line after finishing the message, and the message is mailed. The next time Moe presses the Enter key, or when he logs in, his screen will display:

```
You have new mail.
```

Then, Moe invokes his mail utility of choice to see what mail has arrived.

A setup file called **.mailrc** contains the mail characteristics for each user. For more information on the **.mailrc** file, see the **mailx**(1) manual page. For detailed information on using mail facilities, see *Using the DG/UX™ System* and the **mail**(1), **mailx**(1), and **xdgmail**(1X) manual pages.

Boot Time Responsibilities

When a system boots, it loads and executes its kernel, the program that provides operating system services. After initializing some of its internal functions and, to some extent, the hardware, the kernel starts the **init**(1M) process to start various system services not provided in the kernel itself.

Depending on how you have configured your system, you may have to invoke the **init** command yourself to start the desired services. This is true if your system only boots to run level S or run level 1. At run level S or run level 1, the system provides some basic services but not all of them. Generally, you get the full complement of services at run level 3, which you reach with this command:

```
# init 3 ↵
```

Installing the DG/UX™ System introduced you to the **init** command and run levels, so you should already know to what run level your system boots. Chapter 3 discusses run levels and the **init** command in more detail.

While the system boots, it produces a variety of messages, log files, and other output that you can review to verify that your system is initialized and functioning correctly. The following sections describe briefly how to review the information produced at boot time.

Monitoring the Boot Process

As booting occurs, a number of messages appear on the system console. These messages start with the loading of the kernel and end when the system has reached the run level to which you have set your system to boot. It is good practice to watch these messages as booting occurs, in case an error message appears. As an alternative, you can review the boot messages later by looking at the file **/etc/log/init.log**.

The first messages result from the loading and initialization of the kernel and reflect the current installed version of the DG/UX system. The first messages also tell

something about your hardware, the amount of physical memory configured in the system and the number of processors making up the CPU.

The kernel then proceeds to configure hardware devices on your system. The kernel recognizes only those devices for which you have included a device driver specification in the system file when you built the kernel. For more information on building kernels, see Chapter 4.

After the kernel has initialized itself, it starts the **init** program to continue system initialization and to start system services. The **init** program performs these functions by executing a series of check scripts and setup scripts, all of which produce their own output. The typical array of output messages on a typical system with networking might look like this:

```
Checking local file systems ...
Current date and time is Tue Sep 22 14:30:24 EDT 1992
Checking system files .....
Enabling automatically pushed STREAMS modules .....
Linking short names for /dev device nodes ...
Loading terminal controllers ....
Starting disk daemons .....
Mounting local file systems ..........
Checking for packages that have not been set up ...
Starting miscellaneous daemons ...
Starting Logical Link Control Services ...
Starting TCP/IP network interfaces .........
Starting system logging daemon ....
Starting NIS services ......
Starting NFS lock services ......

NOTE: Pausing for 15 seconds to allow remote systems to
    reclaim NFS locks.

Starting batch services ....
Starting line printer scheduler ....
Saving ex(1) and vi(1) temporary files ....
Starting NFS services .....
Starting TCP/IP daemons ........
Mounting NFS file systems .....

NOTE: See /etc/log/init.log for a verbose description of the
    system initialization process.
```

In general, it is good practice to look for error messages that may appear in the boot display. In addition, there are several areas where you should pay particular attention.

Checking Local File Systems

The file system checker, **fsck**(1M), runs every time you boot the system. When **fsck** runs at boot or during a change of run levels, its output goes to **/etc/log/fsck.log** or **/etc/log/fast_fsck.log**, discussed in a later section.

The system checks file systems only if an unexpected event such as a system crash or a disk or disk controller hardware failure causes service to the file system to terminate abnormally. In these cases, the system will not allow further access to the file system until you have checked it with **fsck**. If you do not wish to reboot the system to start file system checking, you can use the **sysadm** utility (see Chapter 7), or **sysadm**'s File System menu (see Chapter 9), to start **fsck**. Also see the **fsck**(1M) manual page.

Checking System Files

The **chk.system** script performs several functions, including checking for package setup scripts that have not executed, verifying that the DG/UX parameters file is available, verifying that the **/tmp** directory exists and is usable, and adding additional swap areas if defined.

One important function of this part of the boot process is to check for user profiles that do not have passwords. A user profile without a password allows anyone to log in using that username. If such profiles exist, you should assign passwords to them immediately. The **sysadm** User menu provides operations for managing user profiles.

Initially, the DG/UX system has no passwords for the superuser profiles **root** and **sysadm**. It is important for you to assign passwords to these profiles as soon as you install your system. Leaving these profiles without passwords allows anyone to log in with superuser access.

If you installed the X Window System package, there will also be an **xdm** profile that does not have a password. Although the system warns you of this condition at every boot, you do not need to provide a password for **xdm** (it does not constitute a security hazard).

User profiles are in the **passwd**(4) file, located in **/etc**. The **passwd** file is readable by all. Although the passwords are encrypted, any user on your system can identify profiles that have no password.

Checking for Packages

This part of the boot process checks for software packages that are not set up. This check only includes software packages loaded using the utilities in **sysadm**'s Software menu. Typically, these packages require that you load them onto disk and then set them up before you can use them. There is nothing wrong with having a package that is not yet set up; this check simply serves as a reminder for you. You cannot use a package until you set it up. To set up a package, use the utilities in the Software menu.

If this check finds that you have loaded but not set up a release of the DG/UX system software, it begins the setup process without prompting.

Checking /etc/log

A number of the services started at boot leave logs in the **/etc/log** directory. After every boot, it is good practice to check this directory and review the logs.

The **ls**(1) command provides options that allow you to look more easily for log files created during the last boot. After changing to the **/etc/log** directory, execute this command to list files in the order they were last changed:

```
# ls -ltr )
```

The most recently changed files appear at the bottom of the listing. If you check this directory after every boot, you soon learn to tell if a file contains unusual information simply by looking at the size of the file.

As the system creates boot logs, it renames boot logs left from the previous boot. It renames them by adding the suffix **.old** to the file name. In the process of renaming the previous boot logs, any existing logs with the **.old** suffix are deleted.

Some log files you may want to review are:

init.log This file contains messages that appeared during the boot process. The **chk** and **rc** scripts found in **/usr/sbin/init.d** produce these messages.

fsck.log and **fast_fsck.log**
When run at boot or during a run level change, **fsck** produces these log files. When **fsck** finds that a file system is consistent, it logs an entry like this:

```
/dev/rdsk/disk: No check necessary for /dev/rdsk/disk.
```

where *disk* is the name of the virtual disk containing the file system.

When file systems are corrupt, the **fsck** log contains messages for each error found in the file system. See Chapter 9 for a discussion of **fsck** and its output.

nfsfs.log This file contains the output from the **mount**(1M) command, which indicate whether or not the attempts to mount directories failed or succeeded. When the mount fails, the entry includes the error message. To diagnose network-related failures, see *Managing TCP/IP on the DG/UX™ System* or *Managing ONC™/NFS® and Its Facilities on the DG/UX™ System*.

Checking lost+found

One of the functions of the **fsck** file checker is to locate blocks of data that have become disconnected from their files. If the **fsck** utility cannot reconnect the data to its file, it puts the data in its own file and puts the file in a directory called **lost+found** in the file system's mount point directory. The mount point directory is the directory where you have attached the file system to your system's directory structure.

For example, if you have a file system mounted on mount point directory **/sales/accounts**, **fsck** puts any disconnected blocks in files in **/sales/accounts/lost+found**. The files in this directory have names reflecting where **fsck** found them.

It is good practice to check the **lost+found** directory of a file system after **fsck** has checked it. You could use a simple script like the following to list files in all mounted local **lost+found** directories:

```
#!/bin/sh
/etc/mount | /bin/grep ' type dg/ux ' | /bin/cut -d" " - |\
while read DIR
do
  if test -f $DIR/lost+found/*
  then
    echo Found lost file fragments in $DIR/lost+found:
    /bin/ls -l $DIR/lost+found/*
  fi
done
```

You may also find the **file**(1) command helpful. This command determines the nature of a file by classifying it as English text, data, binary executable code, and so on.

Example: a power outage causes your system to crash. When you reboot the system, **fsck** begins checking file systems. When **fsck** is finished, you look in **fsck.log** (or **fsck_fast.log** for file systems mounted for fast **fsck** checking) and see that file system **/sales/accounts** required checking by **fsck**. The file system is now mounted and accessible. You do this:

```
# cd /sales/accounts/lost+found ↵
# ls -l ↵
total 8
-rw-rw-rw-  1 curly  sales    3273 Sep 23 1991 #177431
# file #177431 ↵
#177431:  English text
```

You see that **fsck** found a piece of a file belonging to user **curly**. The **file** command classifies the contents as English text. Now you can tell user **curly** that one of his files was damaged, and, using the fragment from the **lost+found** directory as a clue, he can set about determining what file was damaged so that he can repair it or have it restored from backup.

Depending on how you configure your system, rebooting after a power outage or other failure may occur without operator intervention. Thus, file system checking could occur without your ever realizing it. Therefore it is good practice to check the **lost+found** directories periodically to be sure there are no file fragments there. The absence of a **lost+found** directory simply means that **fsck** has never needed to create one.

Day-to-Day Responsibilities

In addition to checking on the state of your system every time you boot, there are tasks you should perform more frequently.

Monitoring System Health: /var/adm/messages

A number of system facilities produce messages describing errors, abnormal conditions, routine checkpoints, and a number of other events. These facilities use

the system error logger daemon, **syslogd**(8), to direct the messages to appropriate destinations, such as the system console, a user terminal, and files such as **/var/adm/messages**.

It is good practice to review **/var/adm/messages** occasionally to verify that your system is operating as expected. Some messages may indicate conditions requiring your attention. To change how the system handles these messages, see Chapter 16 and the **syslog.conf**(5) manual page. Also, see Chapter 16 for information on generating log reports.

If you run the **/usr/sbin/newsyslog** script periodically, **/var/adm** may also contain old message files named **messages.0**, **messages.1**, **messages.2**, and **messages.3**. The **newsyslog** script is one of the jobs in the prototype **root crontab** file. If you install the jobs in this prototype file, **newsyslog** will run once a week, renaming the various **messages** files so that you can more easily see how old they are. For more information on the **cron** utility and the prototype **root** jobs, see "Automating Job Execution."

Managing Physical Disks

You may administer an entire physical disk with operations that software format, configure and deconfigure, register and deregister, copy contents to another physical disk, map faulty blocks to known good ones, and convert among data storage management formats, among others. To improve your system's performance, you may reconfigure your physical disk resources without taking your system down. You may use both the **sysadm** operations accessible through the `Device -> Disk -> Physical` menu and the **admpdisk** command to perform these operations. See Chapter 7 for information on the former, and the **admpdisk**(1M) manual page for the latter.

Monitoring Disk Space

It is important to keep track of the size of your file systems. When a file system becomes 80% full, I/O performance begins to decline. When a file system becomes 90% full, operations requiring more space (such as creating new files or directories or increasing the size of a file) will fail. When a file system is 90% full, only the superuser can perform an operation that involves allocating more space in the file system. Any non-superuser process that tries to create a file or write to a file in a 90%-full file system will fail with an error such as `no space left on device`.

The **sysadm** operation File System -> File Information -> Disk Use shows the number of blocks used and free and the number of inodes (file slots) used and free.

The operation File System -> File Information -> Find searches for files based on a variety of criteria. For example, you can search for files of a given size, or you can search for files modified during the last week.

You can also use the **df**(1M) command to display the space allocation information for a file system. Following is a sample **df** display.

```
# df -lt / /usr ↵
/     (/dev/dsk/root      ):     8199 blocks    4869 files
```

```
                    total:    40000 blocks    5760 files
/usr   (/dev/dsk/usr     ):     1415 blocks   25895 files
                    total:   240000 blocks   34560 files
```

The display shows the mount point directory name, the pathname of the virtual disk device, the number of free (unallocated) 512-byte data blocks and file slots, and the total number of data blocks and file slots.

To change the size of a file system, see the File System Management menu of **sysadm**, described in Chapter 9.

To display the size of any directory (including any subdirectories), use the **du**(1) command. The **du** command returns the size in 512-byte blocks.

In the / file system, there are several directories and files that you should check regularly to make sure they are not growing excessively. What constitutes excessive growth depends on how much free space you need in your file systems. The following sections describe these files and directories in more detail.

Cleaning out Temporary File Directories

First among these directories are the system temporary file directories located in the / file system, **/tmp** and **/var/tmp**. Various programs use these directories to store temporary work files. Some applications normally do not remove their temporary files, and many applications leave temporary files behind if they terminate abnormally.

The best way to keep temporary file directories clean is by running a janitor-type job regularly using the **cron** utility. The prototype **crontab** file for **root** includes jobs that clean out **/tmp** and **/var/tmp** periodically. For more information on the **cron** utility and the prototype **root** jobs, see "Automating Job Execution" later in this chapter.

Cleaning Up Log Files

Table 2–1 lists each DG/UX log file, its purpose, whether the controlling program appends or truncates it, and a method for cleanup.

According to Table 2–1, if the "Cleanup" column recommends "special" steps, see a subsequent section in this chapter for the appropriate cleanup. If the column entry is "truncate," use this command:

```
# cat /dev/null > filename
```

Instead of truncating the file completely, however, to preserve the newer entries in a log file, deleting the older entries, you may use this command:

```
# tail filename > /tmp/filename; cat /tmp/filename > filename
```

Table 2–1 DG/UX Log Files

Log File Location	Log's Purpose	Usage	Cleanup
/etc Directory Logs			
/etc/log/fsck.log	fsck	truncated	truncate
/etc/log/fast_fsck.log	fast recovery fsck	truncated	none
/etc/log/filesave.log	filesave	appended	truncate
/etc/log/init.log	init	truncated	none
/etc/log/netinit.log	netinit	truncated	truncate
/etc/log/nfsfs.log	nfsfs mount requests	truncated	truncate
/etc/log/preserve.rclock	preserve start time	truncated	truncate
/etc/lp/logs/lpNet	lpNet start	appended	truncate
/etc/lp/logs/lpsched	lpsched start/stop	appended	truncate
/etc/lp/logs/requests	lp request	appended	truncate
/etc/wtmp	user data	appended	See later section on truncating /etc/wtmp file
/etc/utmp	user data	truncated	none
/var Directory Logs			
/var/adm/acct/nite/fd2log	accounting	truncated	none
/var/adm/acct/nite/wmtp.*MMDD*	accounting data	truncated	none
/var/adm/acct/nite/wmtperror*MMDD*	accounting errors	truncated	none
/var/adm/log/backup.log	sysadm dump2	appended	truncate
/var/adm/log/idc.log	ldc compiler	appended	truncate
/var/adm/log/sysadm.log	sysadm	appended	truncate
/var/adm/messages	system messages maintained by syslog	appended	truncate
/var/adm/sa/*	sar data	See **sar(1)** manual page	
/var/adm/spellhist	data	appended	truncate

Log File Location	Log's Purpose	Usage	Cleanup
/var/adm/sulog	switch user	appended	truncate
Note: For all log files located in **/var/adm**, see the later section on "Cleaning Up the /var/adm Directory" for more information.			
/var/cron/log	cron start	appended	truncate
/var/lp/logs/lpNet	lpNet	appended	truncate
/var/lp/logs/lpsched	lpsched	appended	truncate
/var/lp/logs/requests	lp	appended	see later section on lp print service logs
/var/saf/tcp/log	tcp listen port monitor	appended	see following table note
/var/saf/*pmtag***/log**, where *pmtag* is the name of a port monitor	ttymon port monitor service	appended	see following table note
/var/saf/_log	output from saf process	appended	see following table note
Note: For all log files located in **/var/saf**, see the later section on "Cleaning Up the ttymon Related Log Files" and Chapter 10 for more information.			
/var/setup.d/log/dgux.root	setup of root	appended	truncate
/var/setup.d/log/dgux.usr	setup of usr	appended	truncate
/var/setup.d/log/nfs.root	setup nfs root	appended	truncate
/var/set.d/log/onc.root	setup onc root	appended	truncate
/var/setup.d/log/tcpip.root	setup tcpip root	appended	truncate
/var/setup.d/log/tcpip.usr	setup tcpip usr	appended	truncate
/var/spool/uucp/.Log/uucico/system	uucico	appended	truncate
/var/spool/uucp/.Log/uuxqt/system	uuxqt	appended	truncate
Note: For all log files located in **/var/spool**, see the later section on "Cleaning Up the /var/spool Directory" for more information.			

Some of these log files take care of their own cleanup. Others you have to clean up by truncating them. The following sections describe files that require special attention.

Truncating the /etc/wtmp File

You should reduce the size of the **/etc/wtmp** file occasionally. The system logs accounting and user login information to this file on a continual basis, causing the

file to grow. When the file becomes too large, you may choose either to remove it, replacing it with an empty **wtmp** file, or to reduce it, leaving only some of the more recent entries. If you use accounting on your system, you should note that removing or reducing this file removes information required by the accounting system to charge for connect time. See Chapter 14 for more information on the accounting system.

If you remove the oversized **wtmp** file, remember to replace it with an empty file. If you choose to reduce the size of the file, you should note that this file is a data file (not a text file) made up of 64-character entries, and you should not edit it with a text editor. Instead, use the **tail**(1) command to extract entries from the end of the file, the goal being to replace the **wtmp** file with these final entries. On the **tail** command line, be careful to specify a number of characters that is divisible by 64, the size of a file entry. If any entry in the new **wtmp** file is incomplete, some commands may fail. For example, to reduce **wtmp**, leaving only the last 3200 characters, issue these command lines:

```
# tail  -3200c  /etc/wtmp  >  /tmp/wtmp ↵
# mv  /tmp/wtmp  /etc/wtmp ↵
```

For more information, see the **wtmp**(4) manual page.

Cleaning Up the /var/adm Directory

The **/var/adm** directory contains logs of various kinds that you should check occasionally. For example, if you use the system activity monitor, **sar**(1) or **nsar**(1), you need to make sure that **sar** output logs in **/var/adm/sa** do not take up too much space.

If you use the accounting system, you need to check the **/var/adm** and **/var/adm/acct** directories occasionally to make sure they are not growing out of bounds. For more information on these files and the accounting system, see Chapter 14.

If users on your system use the **spell**(1) utility, you should occasionally check **/var/adm/spellhist** to make sure it is not growing excessively. The **spellhist** file contains words not recognized by **spell**. If your users have no use for the **spellhist** file, you may simply delete it; however, if they like to track words that appear there, you can at least reduce the size of the file by using **sort**(1) to sort the file and remove duplicate entries. The following example demonstrates:

```
# cd /var/adm ↵
# sort -u spellhist > /tmp/spell.tmp ↵
# mv /tmp/spell.tmp spellhist ↵
```

Cleaning Up LP Print Service Logs

The LP print service has several logs that require occasional trimming. These logs, located in **/var/lp/logs**, are **lpNet**, **lpsched**, and **requests**. The LP print service's prototype **crontab** file, **/admin/crontabs/lp.proto**, contains jobs to prevent these logs from growing without limit. To use them on your system, execute **crontab –e** to

edit the superuser's **crontab** file. Then add the jobs from the **lp.proto** file. For more information on **cron**, see "Automating Job Execution." For more information on the LP print service logs and the **lp.proto cron** jobs, see Chapter 12.

Cleaning Up the ttymon Related Log Files

When cleaning up log files related to **ttymon**, be careful not to remove or re–create these files while either **ttymon** or **sac** is running. The **sac** log file is **/var/saf/_log**, and **ttymon** log files are stored in **/var/saf/***pmtag***/log**, where *pmtag* is the name of the port monitor.

If you are not interesting in saving the data from these log files, and decide to truncate, use these commands:

For **sac**:

```
# cd /var/saf/;>log
```

For **ttymon**:

```
# cd /var/saf/pmtag;>log
```

If you want to preserve some of the latter data in the log files, use these commands instead:

For **sac**:

```
# tail /var/saf/_log >/tmp/saflog; cat /tmp/saflog >/var/saf/_log
```

For **ttymon**:

```
# tail /var/saf/pmtag/log >/tmp/pmtaglog; cat/tmp/pmtaglog \
>/var/saf/pmtag/log
```

The **>** character is a redirection operator.

Cleaning Up the /var/spool Directory

Occasionally, you should check the **/var/spool** directory, which contains directories supporting a number of system services. For example, this directory contains the files and directories supporting the printer (LP) services, described in Chapter 12. An interrupted LP command could abandon a file in the LP requests directory, for example. If such an accident happens often enough, you could waste valuable disk space. When you find lost print jobs, you can delete them or save them for the user who submitted the print job.

Making Backups of File Systems

Another important responsibility of the system manager is backing up and restoring files. If you have an operations staff at your site, you may not be the person performing the backups, but you may still want to acquaint yourself with the operations involved.

Typically, you start by backing up your entire system at the beginning of the month. Thereafter, you perform backups at the end of every workday except the last day of the week, backing up only those files that changed during that day. At the end of the last day of the week, you back up files that changed during the entire previous week. When the next month starts, you repeat the cycle.

A complete discussion of backing up and restoring file systems is beyond the scope of this chapter. See Chapter 9 for more information.

Maintaining and Verifying System Security

The DG/UX system provides a number of security features. For environments where you require a greater degree of security, however, there are also the Trusted DG/UX systems, which provide B1 and C2 levels of security. For more information on the Trusted DG/UX systems, contact your Data General representative.

In general, you may find the following suggestions helpful for maintaining a secure system.

- Set the access permissions to directories and files to allow only the necessary permissions for owner, group, and others.
- All logins should have passwords. Advise users to change passwords regularly. Password aging, discussed in Chapter 13, is a feature that can force users to change passwords regularly. Advise users not to pick obvious passwords. Users should avoid passwords that common "password-cracker" programs may guess, such as proper names and any word appearing in the dictionary. In addition, the **passwd**(1) command, used to set passwords, imposes other restrictions.
- All port services, whether served through a terminal or a modem, should run **login** or another service that requires password validation before granting access to the system.
- Users who make frequent use of the **su** command can compromise the security of your system by accessing files belonging to other users without the other users' knowledge. The more people who know a given login and password, the less secure access is to the system. For this reason, a log is kept on the use of the command. Check the file **/usr/adm/sulog** to monitor use of the **su** command.
- Login directories, personal configuration files, such as **.profile**, **.login**, and **.cshrc**, and files in **/usr/sbin**, **/usr/bin**, and **/etc** should not be writable by others.
- Encrypt sensitive data files. The **crypt**(1) command and the encryption capabilities of the editors (**ed** and **vi**) provide protection for sensitive information. Encryption/decryption capabilities are available as a separate product with only U.S. releases of the DG/UX system. Contact your Data General representative for more information.
- Do not leave a logged-in terminal unattended, especially if you are logged in as **sysadm** or **root**.

- To check your file systems for files that may indicate security breaches, use the **sysadm** operation `File System -> File Information -> Check`. The operation looks for two kinds of files:

  Device files outside of /dev
  Device files, normally located in the **/dev** directory, provide access to peripheral devices on your system. The Check operation looks for device files in directories other than **/dev** because such files may provide unauthorized access to the data on a device.

  Setuid executables owned by the superuser
  Executables (programs and scripts) that have the setuid bit set will run under the username of the program's owner rather than with the username of the invoking user. If such an executable is owned by the superuser (the **root** or **sysadm** profile), the program will run as a superuser process, regardless of who invokes it. Depending on the nature of the program, it may give an unauthorized user access to sensitive data or applications.

  See Chapter 9 for more information on the Check operation.

- Do not put the current directory, represented by a dot (.), on any superuser search path. In user search paths, the current directory should always appear last, if at all. Placing the current directory on your path may cause you to execute inadvertently an nonsecure script or program that has the same name as a common command.

Monitoring the Mail System

By default, the DG/UX system uses **/var/mail** as the mail directory for use by the **mailx**(1) program. It is good practice to check **/var/mail** occasionally to look for:

Oversized mail files
Simply enough, some users never delete mail. If you find that your **/var** directory is getting too full, check the **mail** directory to see if there are any excessively large files.

Misaddressed mail
Typographical errors when sending mail can result in mail messages that sit indefinitely in files that no one reads. Depending on how your site is configured, your mail system will probably return mail to users if they accidentally send it to a nonexistent user. Nevertheless, you may want to inspect the listing of mail files in **/var/mail** occasionally to make sure that only valid files, named for real users, exist.

Some sites require users to use the **mailx** command instead of the **mail** command because the **mail** command cannot resolve network mailing addresses the way **mailx** can. Accidentally omitting the the **x** from **mailx** may thus result in a mail message being delivered to the local mail directory rather than to a remote system where it belongs.

Superuser or postmaster mail
Some system services alert the administrator to urgent or abnormal events

by sending mail to the **root** profile. Your **root** mail file may also receive messages addressed to **postmaster**, **sysadm**, or another administrative title at your site.

If you do as many administrators do and log in as a normal user, using **su** to become the superuser as needed, you may not see the **root** mail notification message that would appear at login. Therefore, it is good practice to remember to check for **root** mail occasionally, or even to add a line to one of your personal configuration files (such as **.profile**, **.login**, or **.cshrc**) that checks for **root** mail when you log in.

Automating Job Execution

As system administrator, you will find it helpful to be able to schedule jobs to run on a regular basis. For example, you may have a script that checks disk free space and file system security. With the **cron** facility, you can schedule this script to run once a week or every night, for instance. The DG/UX system also offers the **at** and **batch** facilities, which run jobs at low priority or at any time that you specify. The following sections elaborate.

For further information on topics covered in this section, see manual pages for **cron**(1M), **crontab**(1), **at**(1), **atq**(1), **atrm**(1), and **batch**(1).

Scheduling Periodic Jobs with cron(1M)

As system administrator, you will find that the **cron** utility is one of your handiest tools. With **cron** you can schedule a job to run every five minutes, twice a week, or once a year, for example. Many subsystems of the DG/UX system, such as the LP print service and the accounting subsystems, include **cron** jobs for tasks such as collecting data and maintaining logs. You and other users on the system may also submit your own **cron** jobs.

Setting up a **cron** job involves executing the **crontab** command to start an editing session for your **crontab** file. Your **crontab** file contains a one line entry for each scheduled **cron** job. An entry comprises two kinds of information: the frequency of the job and the command line to be run. When you finish editing the file and exit the editor, **crontab** submits your new **crontab** file to **cron**. At the appointed times, **cron** runs the jobs, mailing you the output if any. Note that before other users can create a **crontab** file, the administrator must add their name to the **/etc/cron.d/cron.allow** file.

The following steps detail the procedure for scheduling system administration **cron** jobs as superuser.

1. Issue this command:

   ```
   # crontab -e ↵
   ```

   If there is no **crontab** file for the superuser on your system, the **crontab** command will create an empty one; otherwise, the command opens an editing session with the existing **crontab** file. The **crontab** command invokes the **ed**(1) editor unless your EDITOR environment variable specifies another.

2. Using the editor, insert the desired entries. An entry has the following format:

   ```
   minute  hour  day-of-month  month  day-of-week
   ```

   where:
   minute is in the range 0 – 59.
   hour is in the range 0 – 23.
   day-of-month is in the range 1 – 31.
   month is in the range 1 – 12.
   day-of-week is in the range 0 – 6, with 0 = Sunday.

   When specifying multiple values in a field, separate the values with a comma (,). In any field you may use the asterisk (*) to represent all possible values. The following sections provide examples of and suggestions for **cron** jobs.

   NOTE: If you change the time or timezone or if either daylight savings time (DST) or standard time (ST) begins, stop and restart **cron.**

3. Save the file and exit from the editor. The new **cron** jobs will run at the appointed times.

   NOTE: The **crontab** command will not submit the jobs if the editor returns an exit code other than 0. See the man page or other documentation for your editor. Some editors, such as the DG/UX editors **ed** and **vi**, return a non-zero exit code if you remove all lines from the file. In this case, **crontab** does not submit the file to **cron**, and your previous table of **cron** jobs remains in effect. The proper way to remove all **cron** jobs is to invoke **crontab** with the **-r** option.

   You know that **crontab** has submitted your edited file of jobs when it returns this message:

   ```
   warning: commands will be executed using /usr/bin/sh
   ```

   This warning does not indicate a problem; rather, it serves to remind you that the command lines in your job entries should adhere to Bourne shell syntax. See **sh**(1) for more information on the Bourne shell.

Here are some suggestions for scheduling **cron** jobs:

- Try to schedule jobs for off-peak hours, particularly if the job is expensive in terms of resources like CPU time or disk access. Running during off-peak hours, the job is less likely to inconvenience other users. Example 2 demonstrates this point.
- Schedule jobs to run at odd minutes or hours to avoid coinciding with other jobs. If you and other users typically run your jobs only on the hour or half hour, you may find a noticeable degradation in system performance at these times if multiple jobs run simultaneously. Example 3 demonstrates this point.
- Minimize security risk by specifying complete pathnames of commands and by executing only commands residing in secure directories. A **cron** job executes as the user who submitted the job; therefore, it is particularly important that jobs to be run as superuser execute only secure commands. Scripts especially should have permissions set to prevent tampering.

Examples

1. The following entry schedules a **wall**(1M) message to send out a reminder every Tuesday morning at 9:00:

```
0 9 * * 2 /bin/echo /bin/echo "Meeting at 9:30!" | /etc/wall
```

2. The following entry schedules a file system security checking job to run on Monday, Wednesday, and Friday mornings at 2:15:

```
15 2 * * 1,3,5 /usr/sbin/admfsinfo -lq -o check
```

   The **cron** facility will use **mail** to send you the output from the command line.

3. The following entry schedules a script to run every hour, Monday through Friday, at 10-minute intervals starting at 3 minutes past the hour. The script appends output to a file:

```
3,13,23,33,43,53 * * * 1-5 /admin/getCPUload >> /admin/load.log
```

4. The following entry schedules a message to run at 2:07 in the afternoon on June 20:

```
7 14 20 6 * /bin/echo /bin/echo "Happy Summer Solstice!" | /etc/wall
```

5. The following entry schedules the **ntpdate** function to poll a clock server on your LAN to reset your host time every eighth minute of each hour.

```
8 * * * * /usr/bin/ntpdate commtg3 brewery dg-rtp >/dev/null 2>&1
```

Prototype Jobs for System Administration

Although the DG/UX system does not by default have any scheduled **cron** jobs, it does provide some prototype jobs that you may adapt or use as shipped. These jobs are in the following files in **/admin/crontabs**:

root.proto This file contains two kinds of jobs: those helpful on systems running accounting, and those helpful for systems in general. The file contains:

```
0 4 * * * /bin/su - adm -c "/usr/lib/acct/runacct 2> \
  /usr/adm/acct/nite/fd2log"
5 * * * * /bin/su - adm -c "/usr/lib/acct/ckpacct"
15 5 1 * * /bin/su - adm -c /usr/lib/acct/monacct
0 2 * * * /usr/lib/acct/dodisk
0 3 * * 2-5 /bin/find /tmp /var/tmp -mount -atime +3 \
  -type f ! -name '[XM][0-9]*' -exec rm {} \;
15 3 * * 6 /bin/find /var/adm/log -name \
  'sysadm.log.[0-9][0-9][0-9]' -atime +7 -exec rm {} \;
50 9,3 * * * /bin/find /var/spool/cron/atjobs
  -name '.nfs*' -atime +1 -exec rm {} \;
5 4 * * 6 /usr/lib/newsyslog >/dev/null 2>&1
```

Some lines have been broken for readability. The first four jobs perform accounting functions. If your system does not use accounting (covered in Chapter 14), you do not need to schedule these jobs. The fifth through

the eighth jobs clean up temporary file directories and remove outdated logs and some other files.

lp.proto This file contains jobs to help maintain the LP system. You should schedule these jobs on any system that uses printers. The file contains these jobs:

```
13 3 * * * cd /var/lp/logs; if [ -f requests ]; then \
  /bin/mv requests xyzzy; /bin/cp xyzzy requests; \
  >xyzzy; /usr/lbin/agefile -c2 requests; /bin/mv \
  xyzzy requests; fi
15 3 * * 0 /usr/lbin/agefile -c4 /var/lp/logs/lpsched
17 3 * * 0 /usr/lbin/agefile -c4 /var/lp/logs/lpNet
```

The job beginning on the first line has been broken over multiple lines for readability. You should run these jobs as **lp**. To schedule these jobs on your system, first give **lp** permission to run **cron** jobs by adding an **lp** entry to **/etc/cron.d/cron.allow**. Then execute **su lp** command to become **lp** before executing **crontab –e** to schedule the desired jobs. See Chapter 12 for more information on the LP print service.

uucp.proto This file contains jobs for maintaining the UUCP file transfer and remote command execution facility. The prototype jobs are:

```
39,9 *    * * *  /etc/uucp/uudemon.hour  > /dev/null
10  *    * * *  /etc/uucp/uudemon.poll  > /dev/null
45  23   * * *  /etc/uucp/uudemon.cleanup > /dev/null
48  10,14 * * 1-5 /etc/uucp/uudemon.admin  > /dev/null
```

If you use UUCP, these jobs should run as **nuucp**, the login name intended for UUCP administration. To schedule these jobs on your system, first give **nuucp** permission to run **cron** jobs by adding an **nuucp** entry to **/etc/cron.d/cron.allow**. Then execute **su nuucp** command to become **nuucp** before executing **crontab –e** to schedule the desired jobs. See *Using Modems and UUCP on the DG/UX™ System* for more information on UUCP.

Maintaining the cron Log

The **cron** utility logs a history of its activity to the **/var/cron/log** file. You should periodically truncate the **/var/cron/log** file to prevent it from using too much disk space.

1. Stop **cron** by executing the following command line:

```
# /usr/sbin/init.d/rc.cron stop ↵
```

2. Use the following command line to remove all but the last 100 lines of the log file:

```
# tail -100 /var/cron/log > /tmp/log; mv /tmp/log /var/cron/log ↵
```

3. Restart **cron** with the following command line:

```
# /usr/sbin/init.d/rc.cron start ↵
```

Submitting Jobs for Delayed Execution with at(1)

To run a job at a specified time on a specified date, use the **at**(1) command. This command is useful for running jobs during off hours or at any time when you may not be available or may forget to run the job yourself.

For example, to run the script **/admin/check_disks** at 3:00 in the morning on January 24, issue this command line at the shell prompt:

```
# echo /admin/check_disks | at 3:00 january 24 ↵
```

To broadcast a reminder about a meeting one hour from now, issue this command line:

```
# echo "echo Meeting at 2pm | wall" | at now + 1 hour ↵
```

To reboot the system at 11:00 tomorrow night, use this command line:

```
# echo init 6 | at 23:00 tomorrow ↵
```

NOTE: The method of rebooting shown above is not recommended for active systems. See the section on shutting down the system in Chapter 3 for more information.

That **at** command accepts date and time specifications in a variety of formats. For more information, see the **at**(1) manual page. The **at** command submits jobs to the **a** queue managed by **cron**. For more information on queues, see the **cron**(1M) manual page.

Submitting Low-Priority Jobs with batch(1)

There are times when you want to execute a job immediately but hesitate to do so because CPU time is in demand. At these times, the **batch** command is useful because it submits the job to a low-priority queue intended to minimize impact on system performance. As with the **at** and **cron** utilities, the **batch** utility mails you any output from the job.

For example, to submit the script **/admin/check_disks** at a low priority, use the following command line:

```
# echo /admin/check_disks | batch ↵
```

The **batch** command submits the job to the **b** queue managed by **cron**. For more information, see the **batch**(1) and **cron**(1M) manual pages.

End of Chapter

Chapter 3
Operating the DG/UX System

This chapter shows you how to perform operations such as starting the system, shutting it down, recovering from trouble, collecting error messages, and dumping the system memory image and kernel file for analysis by Data General. The chapter also explains how the **init**(1M) command uses **rc** scripts to set run levels and thus provide system services.

Operational Terms

Read the following definitions before beginning the procedures in this chapter:

init The **init**(1M) program creates all system processes based on entries in the file **/etc/inittab**. **Init** is invoked in two ways: inside the DG/UX system as the last step in the boot procedure, and from the command line with a run level as argument. When invoked during the boot procedure, its first function is normally to start a single superuser shell for the system console.

run level Run levels, also known as run states or run modes, provide varying degrees of service on the system. These services include network-related capabilities, accounting, **cron** batch job scheduling, line printer services, and so on. Typically, run level S (for single-user mode) provides no services, and run levels 1 through 3 provide increasing levels of system functionality. By default, DG/UX provides full multiuser and network capabilities at run level 3.

rc scripts The run command scripts are executed at every boot and every time you change run levels. At boot time or whenever run levels change, the **init** program executes scripts in a directory set up specifically for the intended run level. These scripts are the "run command," or **rc** scripts. They kill or start system services as directed by prefixes that consist of either a **K** (kill) or an **S** (start) followed by an ID number. Upon entering a run level (via the **init** command), all **rc** scripts designated in that run level are executed. Execution means that services are killed in order from highest ID number to lowest ID number. Next, processes are started in order from lowest ID number to highest ID number. The result is that only certain **rc** scripts, those that are started with the **S** switch, are active in any run level.

SAF The SAF (Service Access Facility) manages ports, setting terminal type, mode, speed, and line discipline characteristics for the port. SAF can also start service programs for ports. An important function of SAF is to control user terminal lines, starting the **login** service for users who need to log in.

SCM> The System Control Monitor prompt, displayed when no operating system is running on your computer. This prompt comes from your computer hardware. From the SCM, you use the **b** command to boot (begin execution) of your DG/UX kernel program, thereby starting the operating system. At the SCM prompt, type **h** for help, or see *Testing and Troubleshooting AViiON® Computers: AV/Alert and the AViiON® System Control Monitor* for complete information.

System Services

There are a number of system services that the DG/UX system starts when it boots. The services become active at various run levels. Run levels are discussed in the next section. These services (not necessarily in order) include:

Package Setup Check
This service checks to see if any software packages on your system are not set up.

Password Check
This service looks for user profiles that do not have passwords.

File System Checker
This service verifies file system integrity.

Local and Remote File Systems
This service makes local and remote (network) file systems available.

Editor Preservation
This service restores editing sessions that may have terminated abnormally.

Batch Job Services
This service manages batch jobs and jobs that run automatically on a regular basis.

System Error Logger
This service handles messages produced by various other system services.

Terminal Lines
This service provides support for terminals and other ports.

Line Printers
This service provides line printer and print queue support.

Accounting
This service accumulates system statistics for accounting purposes. By default, accounting is turned off.

Miscellaneous Service Daemons
This service starts miscellaneous daemons (background processes) that provide a variety of services.

Network Services
These services manage the network software.

You can also define your own services for invocation at system boot time. These services and how you can add your own are covered later in the chapter.

DG/UX System Run Levels

You can define the run levels to provide whatever services you choose. The services themselves are controlled by the run command scripts (or **rc** scripts) located in **/usr/sbin/init.d**. The mechanism that ties a given service to a run level is the **init** program and the link files located in the **/etc/rc***n***.d** directories (where *n* is **S**, **i**, or one of the numerals **0** through **6**).

Table 3–1 shows the default run levels for the DG/UX system.

Table 3–1 DG/UX Run Levels

Run Level	Description
0	Halts the system.
i	Installation mode starts the **installman**(1M) command, which leads you through installation of the DG/UX system. At this run level, local file systems and disk services are available.
S	Single-user is a low-level run mode that is the default level the system enters upon booting. The only process running is **init**. Only the / (root) and **/usr** file systems are available.
1	Administrative mode is used to install and remove software utilities, run file system backups and restores, and check file system integrity. All local file systems are mounted. Only processes associated with the system console may run.
2	All local file systems are mounted. Local users can log in at terminals and use local facilities, and some out-bound network services are available. Outside systems cannot contact this system over the network. ONC/NFS services are not available.
3	The normal running mode of the DG/UX system. Complete multiuser and network services are available. ONC/NFS services are available. On workstations, the X Display Manager (**xdm**(1X)) is running.
4	User-defined run level. By default, this run level is the same as run level 3.
5	Halts the system.
6	Halts and reboots the system.
a, b, c	Pseudo run levels. These can be specified without changing a run level. Typing **init a**, for instance, invokes those entries in **inittab** that have an **a** in the **level** field. See **init**(1M).

See "Expert Run Level Information" at the end of the chapter for more information.

Default Multiuser Conditions

See Table 3–2 for a complete list of the scripts that run when you go to multiuser states 2 or 3.

When you bring up the DG/UX system in multiuser state, the following things happen:

- Local file systems are mounted in run level 2 (as they are in run level 1).
- Remote file systems are mounted in run level 3.
- The error daemon, the batch job scheduler, various disk-related services, the network status monitor, and the network lock daemon are started.
- The LP system and UUCP are ready to use. The SAF (Service Access Facility) monitors ports, providing whatever services you have configured it to provide. SAF's most notable service is monitoring user terminal lines and starting the **login** program for users who want to log in.
- If used, TCP/IP transmissions work outward in run level 2, and work in both directions in run levels 3 and 4.

Setting Run Levels

This section describes how the **init**(1M) command, the **inittab**(4) file, and the **rc** (run command) scripts define a run level and determine what processes and services are available on your system.

Consider the case where you want to make more service available to a system currently running at single-user mode. Here is what you do:

1. Enter the **init 2** command to change run levels upward from single-user mode S.
2. The **init 2** command causes the **init** program to read the **inittab** file looking for all entries containing the number 2 in the *level* field. **Init** executes all lines that have 2 in the *level* field.
3. The **init** program executes all **rc** scripts associated with run level 2. These scripts perform tasks such as turning on accounting, starting the LP scheduler, and starting various daemons. Run level 2 is therefore defined as all those script-started processes running as a result of the **init 2** command. Output from the **rc** scripts appears in **/etc/log/init.log**.

Run Command Scripts Per Run Level

You can read the **rc** scripts to see exactly what they do. We recommend that you do not modify these scripts. You can add your own if needed. See *Porting and Developing Applications on the DG/UX™ System* for directions.

For systems with the TCP/IP and ONC/NFS packages (in addition to the DG/UX system) loaded and set up, Table 3–2 shows which scripts are started per run level.

Note the cumulative effect: the higher the run level, the more processes are running. Blanks indicate that a script is not running.

Table 3–2 RC Scripts Per Run Level

i	S	0	1	2	3	4	5	6
	rc.ups	rc.ups	rc.ups	rc.ups	rc.ups	rc.ups		
			rc.tcload	rc.tcload	rc.tcload	rc.tcload		
rc.update			rc.update	rc.update	rc.update	rc.update		
rc.localfs			rc.localfs	rc.localfs	rc.localfs	rc.localfs		
			rc.sync	rc.sync	rc.sync	rc.sync		
			rc.lan	rc.lan	rc.lan	rc.lan		
			rc.setup	rc.setup	rc.setup	rc.setup		
			rc.daemon	rc.daemon	rc.daemon	rc.daemon		
rc.install								
				rc.usrproc	rc.usrproc	rc.usrproc		
				rc.llc	rc.llc	rc.llc		
				rc.syslogd	rc.syslogd	rc.syslogd		
				rc.dgserv	rc.dgserv	rc.dgserv		
				rc.account	rc.account	rc.account		
				rc.cron	rc.cron	rc.cron		
				rc.lpsched	rc.lpsched	rc.lpsched		
				rc.preserve	rc.preserve	rc.preserve		
					rc.failover	rc.failover		
		rc.halt					rc.halt	
								rc.reboot
				rc.tcpip-port	rc.tcpip-port	rc.tcpip-port		
					rc.tcpip-serv	rc.tcpip-serv		
				rc.ypserv	rc.ypserv	rc.ypserv		
				rc.nfslockd	rc.nfslockd	rc.nfslockd		
					rc.nfsserv	rc.nfsserv		
					rc.nfsfs	rc.nfsfs		

The following section defines what the scripts in **/usr/sbin/init.d** do, and shows at which run levels they are in effect. See "Expert Run Level Information" for a table showing the kill/start mechanism for all scripts active at all run levels.

RC Scripts

The RC scripts are located in **/usr/sbin/init.d**.

rc.ups Starts the UPS (uninterruptible power supply) daemon (for systems with the UPS subsystem hardware only). This script runs in single-user mode and in levels 0 through 4.

rc.tcload Loads the SYAC driver code once for run levels 1, 2, 3, and 4.

rc.update Starts various disk-related services in run levels i, 1, 2, 3, and 4. These services include the block I/O daemon (**biod**(1M)) and the write-verify service.

rc.localfs Mounts local file systems listed in **/etc/fstab** in run levels i, 1, 2, 3, and 4; unmounts them in all other run levels. A local file system is one of type **dg/ux**, **ramdisk**, **dos**, or **cdrom**. In actuality, the value of the **localfs_ARG** variable, set in **/etc/dgux.params**, determines which file system types to mount.

rc.sync Loads the synchronous controllers used for wide-area network (WAN) communication. This script runs at run levels 1 through 4.

rc.lan Loads the controllers used for local-area network (LAN) communication. This script runs at run levels 1 through 4.

rc.setup Displays packages that have not been set up at initial boot.

rc.daemon Starts miscellaneous daemons.

rc.install Performs installation of the DG/UX system. This script runs only at run level i.

rc.links Create, list, or remove links in the **/etc/rc?.d** directories, where **?** is a regular expression pattern–matching metacharacter. This runs only when you set up the DG/UX system (not during a regular change of run level). You can, if you wish, use **rc.links** to create, list, or remove your own links. This file is a binary executable rather than a shell script.

rc.usrproc Kills all user processes in run levels S, 0, 1, 5, and 6.

rc.llc Starts the llc daemon (**llcd**), which provides logical link control services in run levels 2, 3, and 4.

rc.syslogd Starts the **syslog** error logging program in run levels 2, 3, and 4; kills it in all other run levels.

rc.dgserv Starts DG/UX system services in run levels 2, 3, and 4. This script starts the **dgsvcd** daemon, which provides services for the AV/ALERT facility.

rc.account Starts the **/usr/lib/acct/startup** services and processes in run levels 2, 3, and 4; stops those processes in all other run levels.

rc.cron Starts the **cron** daemon in run levels 2, 3, and 4; kills it in all other run levels.

rc.lpsched Starts the **lpsched** daemon in run levels 2, 3, and 4; kills it in all other run levels.

rc.preserve Invokes the **expreserve** command in run levels 2, 3, and 4 to recover editor files saved during a system crash.

rc.failover Starts **failoverd**(1M) for communicating with another host used for failover disks. This script runs at levels 3 and 4.

rc.halt Halts the processor, taking it to the SCM. This script runs at run levels 0 and 5.

rc.reboot Halts and reboots the system. This script runs only at run level 6.

rc.tcpipport Sets hostname, host ID, network security, and initializes network I/O boards in run levels 2, 3, and 4. These are not set in any other run levels.

rc.tcpipserv In run levels 3 and 4, starts whichever TCP/IP daemons are defined to run on your system. The TCP/IP daemons are **telnetd**, **ftpd**, **tftpd**, **smtpd**, **rlogind**, **rwhod**, **rshd**, and **rexecd**. The **rc.tcpipserv** script kills them in all other run levels.

rc.ypserv Starts the **yp** and **portmap** daemons, and sets the domain name in run levels 2, 3, and 4; kills these in all other run levels.

rc.nfslockd Starts daemons for ONC/NFS file locking, **statd** and **lockd**.

rc.nfsserv Starts the **portmap**, **rwalld**, **mountd**, **ruserd, nfsd**, and **biod** daemons in run levels 3 and 4; kills them in all other run levels.

rc.nfsfs Mounts all local and ONC/NFS file systems listed in **/etc/fstab** in run levels 3 and 4; unmounts them in all other run levels.

Check Scripts

In addition to the run command scripts, the DG/UX system uses several other scripts to set up a properly running environment. These are executed when the system is booted via the **bootwait** action in **/etc/inittab**. Each of these scripts is executed upon the first run level change to levels 1, 2, 3, or 4. For instance, if you boot the system and then go to run level 1, all check scripts are executed. If you then go to run level 2 (without rebooting), then the check scripts are not executed again.

The check scripts are:

chk.date Displays the current system date and allows the administrator to set the correct date. A correct date setting is vital to ensure file creation and modification dates are correct. Also sets time zone based on the **/etc/TIMEZONE** file.

chk.fsck Runs **fsck** on all file systems listed in **/etc/fstab**. The **fsck** program is called with the **-xlp** switch to check file systems in parallel, checking only those file systems that need checking.

chk.system Performs the following system cleanup and initialization routines:

- Initializes the **/etc/ps_data** file.
- Cleans out the **/var/spool/locks** used by the **uucp** program.

- Makes a **/tmp** directory if one doesn't exist.
- Runs the DG/UX setup scripts via the **init** command the first time the system is booted.
- Checks for accounts without passwords.

chk.devlink At the first run level change, this script automatically creates shortened names for devices in the sequence in which it finds them. For example, the first tape device will be device 0, the second will be device 1. These could then be specified as **/dev/rmt/0** and **/dev/rmt/1**. Device short names are taken from the **/etc/devlinktab** file.

chk.strtty This script initializes terminal ports by pushing the required STREAMS modules. The script initializes **duart** and **syac** lines and pseudo-terminals.

Operational Procedures

Operational procedures are necessary to keep the system running on a day-to-day basis.

Starting and Restarting the System

To start the DG/UX system at the SCM prompt, use the **b** (for **boot**) command. If you set the default boot path with the SCM's **f** (for **format**) command, you can boot with this command:

```
SCM> b ↵
```

If the default boot path is not set, you need to specify a boot path. To boot from the first SCSI disk on an AV5220 computer system, for example, use this command line:

```
SCM> b sd(cisc(),0)root:/dgux ↵
```

CAUTION: If your system is part of a dual-initiator configuration and shares a SCSI bus with another system, be careful to specify the correct SCSI adapter SCSI ID in the boot path. If the boot path you use to boot your system specifies the SCSI ID of the SCSI adapter installed in the other system, the boot will fail and the SCSI bus will hang. If the SCSI bus hangs, an attempt by either system to access the SCSI bus will hang the system. When this happens, recover by resetting the hardware and rebooting. See the discussion of failover disks in Chapter 8 for more information.

You can use the **boot** command to load any stand-alone, machine-executable file you choose. A stand-alone, machine-executable file is one that can run directly on the AViiON system hardware without an operating system. As shipped, the DG/UX system's bootable files are:

/dgux A kernel configured for all devices in the standard locations.

/dgux.installer
A hard link pointing to the same file as **/dgux**, above.

/dgux.starter
A hard link pointing to the same file as **/dgux** and **/dgux.installer**, above.

/usr/stand/sysadm
Stand-alone **sysadm**, used for disk management operations when the system on disk is unavailable or for operations on the system disk itself.

In the previous example boot command line, we loaded our standard operating system. If, for instance, you wanted to boot stand-alone **sysadm**, **/usr/stand/sysadm**, located on the first SCSI disk, you would execute it with a command line such as this:

```
SCM> boot sd(cisc(),0)usr:/stand/sysadm ↵
```

You can boot files that reside on file systems built on multiple partitions (to form an aggregation) as long as they are located on one physical disk. Such an aggregation cannot span multiple physical disks. For example, if your **/usr** file system were built on a virtual disk comprised of more than one partition, you could boot **/usr/stand/sysadm** as as long as the partitions were located on the same physical disk.

Booting Alternate Roots and Swap Areas

To boot an alternate virtual disk for **root** and **swap**, include the **–q** option on the boot command line. The command prompts you for the names of the virtual disks to be used in place of **root** and **swap**. For example:

```
SCM> b sd(cisc(),1)root2:/dgux -q ↵

Booting sd(cisc(),0)root:/dgux  -q
Swap disk name (or q to quit)? [swap]  newswap ↵
Root disk name (or q to quit)? [root]  root2 ↵
```

If the kernel cannot find the default virtual disks that you specified (perhaps because of an accidental deletion), the system attempts to find virtual disks named **swap** and **root** on the disk from which you booted. If that fails, the system will attempt to find virtual disks **swap** and **root** on any registered disk. If this fails, the system prompts you again for the names of the **swap** and **root** virtual disks to use.

Alternatively, you may set a default **root** and **swap** virtual disk to be booted from a given physical disk. You may establish default boot settings for use with the **dg_sysctl**(1M) command, which initiates an unattended automatic system reboot in the event of a panic.

Build the virtual disks that you want to use as **root** and **swap**; for example, assume you have **root_production**, **root_test**, and **swap**. If you normally prefer booting **root_production** and **swap**, you could set boot defaults on the physical disk using the **admpdisk**(1M) command, as follows;

```
#admpdisk -o set_defaults -r root_production -s swap 'sd(ncsc(),0)'↵
```

This enables you to set your SCM boot path without specifying the **root** virtual disk or the **–q** option. An example follows:

```
SCM> b sd(ncsc(),0)/dgux -3 ↵
```

To reboot the system using **root_test**, reset the defaults on the physical disk using the following command:

```
# admpdisk -o set_defaults -r root_test 'sd(ncsc(),0)'↵
```

You may reboot without specifying the **root** virtual disk or the **–q** option.

When you boot the system from a physical disk that contains default settings, the bootstrap will attempt to find the named kernel image (**dgux** if it is not specified on the command line) on the virtual disk that is associated with the default **root** virtual disk.

During kernel initialization, the system looks for default settings on the disk from which the system was booted. If it finds settings, it mounts the default **root** virtual disk as the / file system, and begins swapping on the default **swap** virtual disk.

You can remove the **root** and **swap** default virtual disk settings by specifying a null argument in the form of empty double quotes. For example:

```
# admpdisk -o set_defaults -r " " -s " " 'sd(ncsc(),0)'↵
```

To set different boot defaults on different physical disks, maintaining multiple copies of DG/UX on your system, we recommend that you install the DG/UX system on one physical disk using names such as **root_1**, **swap_1**, **usr_1**, and so on. On another physical disk, for another version of the operating system, adopt another naming convention: **root_2**, **swap_2**, **usr_2**, and so on. We strongly urge you to assign unique names to each set of virtual disks to avoid confusion. For example:

```
# admpdisk -o set_defaults -r root_1 -s swap_1 'sd(ncsc(),0)'↵

# admpdisk -o set_defaults -r root_2 -s swap_2 'sd(ncsc(),1)'↵
```

To boot from the first root, use the command line:

```
SCM> b sd(ncsc(),0) ↵
```

To boot from the second root, use the command line:

```
SCM> b sd(ncsc(),1)↵
```

The physical disk has associated with it the virtual disks containing the kernel image, the / file system, and swap space.

NOTE: You cannot assign alternate locations for the **usr** virtual disk since it is mounted after the kernel initializes. If you have multiple **/usr** file systems,

you must set up the proper **/etc/fstab** entries in each of the root file systems you will be booting, referencing the corresponding **/usr** file system to use.

Booting Sequence

If you boot a DG/UX kernel, the boot sequence begins by displaying a message that includes the boot path:

```
Booting sd(cisc(0),0)root:/dgux loading...
        .
        .
INIT: New run level: S

INIT: SINGLE USER MODE
```

Booting brings the system up to the default run level, set in **/etc/inittab**. If you want to boot to a run level other than the default, use a complete **b** command (one that includes the device and boot file specification) and append, as an option, the run level to which you want to boot. For example, to boot from a SCSI disk to run level 2, use a command line like this:

```
SCM> b sd(cisc(),1)root:/dgux -2 ↵
```

You can use the **reboot**(1M) command to halt the system and boot it without placing you in the SCM:

```
# reboot ↵
```

The **reboot** command, without arguments, boots the system using the boot command line used the last time the system booted. Optionally, you may specify a different boot path. Another way to reboot is to use **init** to change run level to **6**, which is the same as executing the **reboot** command. See the **reboot**(1M) manual page for more information.

After a power outage, the automatic boot mechanism reboots your system without operator intervention, bringing it up to the default run level set in **/etc/inittab**. Initially, the default run level is **s** (single-user mode). To change it, edit your **/etc/inittab** file and change this line:

```
def:s:initdefault:
```

so that the second field contains the desired default run level. For example, the following line makes run level 3 the default:

```
def:3:initdefault:
```

Changing Run Levels

You must be superuser to change run levels. You change run levels with the **init**(1M) command. Use this command line to go to administrative mode:

```
# init 1 )
```

Take the system to multiuser mode:

```
# init 2 )
```

Take the system to multiuser mode with network services:

```
# init 3 )
```

Changing run levels causes the **init** command to run various scripts. Many of these scripts write output to **/etc/log/init.log**.

You can also use the **shutdown**(1M) command to take the system down to single-user mode.

Shutting Down the System

As superuser, you can shut down the system from the shell by changing to the root directory and using the **shutdown**(1M) command. The **sysadm** utility does not offer an operation for shutting down the system.

Shutting down the system means taking it to a lower run level. Often, you take the system to run level 1, the administrative state, to perform certain administrative tasks. At other times, you may want to shut down the system so you can halt the processor. In either case, you use the **shutdown** command. In single-user mode, you can use the **halt**(1M) command to stop the processor.

When you shut down the system, system buffers are flushed, open files are closed, user processes and daemons are stopped, file systems are unmounted, and file system superblocks are updated. See "How Run Levels Are Set" later in this chapter for details of what happens as the system comes down through various run levels.

With no options, **shutdown** defaults to run level S, single-user mode.

Shutting Down to Administrative Mode: Run Level 1

Let's assume we're currently in run level 3 and we want to go down to run level 1. In the following example, we'll use the **–i** option to change run levels downward.

Options you can use are:

–y Answers the confirmation query so that shutdown will continue without further user intervention. A default of 60 seconds is allowed between the warning message and the final message. Another 60 seconds is allowed between the final message and the confirmation.

–i1 Go to run level 1, administrative mode.

–g0 Allow a grace period of 0 seconds between the warning message and the final message.

Type the following to change to the root directory and the system down to run level 1:

```
# cd  / ↵
# shutdown  -y  -i1  -g0 ↵

Shutdown started.

The system will be shutdown in 0 seconds.
The system is coming down. Please wait.

INIT: New Run Level: 1
#
```

Now you are in run level 1, administrative mode. Local file systems are the only ones mounted. If you want to shut down to power off, type the **shutdown** command again. You will go to run level S, single-user mode.

Shutting Down to Single-User Mode and Power Off

You can shut down to run level S from any other level. This example shows a shut down from run level 1. Type:

```
# cd  / ↵
# shutdown  -g0 -y ↵
```

After a few moments, you will be in single-user mode. You can change run levels *upward* at this point with the **init**(1M) command, or, you can use the **halt**(1M) command to stop the processor:

```
# halt  -q ↵

CPU HALTED


SCM>
```

You can also halt the system by using **init** to change to run level **0** or **5**. Once at the SCM prompt, you may turn off power to the computer.

Responding to Status Messages

You may be alerted to a class of recoverable operating problems through status codes, which are presented on the system console and are logged to the file **/usr/adm/messages**. Check the file **/usr/release/dgux_5.4R3.00.status.codes** for the status message number and a brief description of the problem. An example follows:

```
IO_EIO_DEVICE_TIMED_OUT                                0x100a18 04005030
```

To find the file containing information on status code `04005030`, you would execute the following commands:

```
# cd /usr/release ↵
# more dgux_5.4R3.00.status.codes ↵
```

Then search for the desired code in the status code file using this command:

```
/5030
```

You then read the indicated file using a command such as **view** or **more**, both of which offer commands for locating the desired text. See the **view**(1) or **more**(1) manual page.

```
IO_EIO_DEVICE_TIMED_OUT                                    0x100a18 04005030
    Device did not respond within timeout period.
```

Most of these messages relate to minor hardware problems, that, once corrected, allow continued normal operation. For example, you could receive a message when attempting to access a file that is located in a disk drive that is turned off. In most cases, the problem is recoverable. However, other conditions may eventually lead to a hang or a panic.

In other cases, the system warns of problems that you should investigate. For example:

```
SCI_EINTR_NFS_SERVER_PROCESS                               0x6c0814 33004024
    An NFS daemon process was terminated by a signal.
```

This message indicates that network services have been suspended. You need to troubleshoot the problem to restore NFS to proper operation.

Using the Watchdog Timer to Detect and Recover from System Hangs

Introduced in DG/UX 5.4R2.10 to support the AV8500, AV9500 and AV/5500 systems, the watchdog timer is a feature that performs an automatic system reset upon detection of either a hardware or software system hang. On AV8500 and AV9500 systems, a reset starts powerup diagnostics, and isolates and deconfigures any faulty field replacable hardware components when the system is rebooted. If AV–Alert (a comprehensive diagnostic support system for AViiON family hardware) is enabled, the faulty hardware will be reported automatically to the Data General Customer Support Center.

DG/UX 5.4R3.00 not only extends support of the watchdog timer feature to the AV4500, AV450, and AV550 systems, but also to all AViiON systems that use the failover monitor **failovermon**(1M). The failover monitor **wdt** should be used in conjunction with DG/UX 5.4R2.10 hardware **wdt** when possible. Refer to Chapter 8 for more information on the failover monitoring process.

Two panic codes may be produced during a hardware or software system hang: 53000060 or 53000061. Panic code 53000060 produces a hardware watchdog timer reset only. Since the current state of the job processors is not saved with a system

reset, a dump of the memory image is not taken. Panic code 53000061 causes the generation of a memory image that can be dumped to tape or disk for analysis.

Enabling the Watchdog Timer and the Failover Monitor

The watchdog timer feature is configured in the kernel by default. Its appearance in the system file follows:

```
# Watchdog Timer
#
# The supported models are:
#
# wdt:  integrated watchdog timer
#
#                              Streams
# Name         Restrictions    Concurrency
# Prefix       Flags           Set
  -------      ------------    -----------
  wdt           o              default
#
```

To enable this feature, ensure that the kernel contains the **wdt()** pseudo device and set up the failover monitor for the software watchdog timer using instructions in Chapter 8.

Setting Parameters for use with the Watchdog Timer

Table 3–3 lists the parameters you may set when using the watchdog timer: the **dg_sysctl** command, **sysadm**, and the SCM.

Table 3–3 Settings for Quick Recovery from System Hangs

Parameter	Value or Path	Where To Set
Autoboot	**auto**	**dg_sysctl** command
Bootpath	**sd(ncsc(),0)root:/dgux -3**	**dg_sysctl** command
Autodump	**auto**	**dg_sysctl** command
Dump device	**ldm_dump(sd(ncsc(),1),sys_dump)**	**dg_sysctl** command
Dump level	**kernel**	**dg_sysctl** command
Power off state	**N/A**	**dg_sysctl** command
Auto–reboot/Boot on Error	**enabled**	SCM
Default boot path	**sd(ncsc(),0)root:/dgux -3**	SCM
Watch dog timer (wtd) configured	`System -> Kernel -> Build`	**sysadm**
Failover monitor set up	`Availability -> Disk Failover -> Alternate Paths -> Add`	**sysadm**

All values set through the **dg_sysctl** command and the SCM must be consistent. For example, if you specified an automatic reboot through the **dg_sysctl** command, you must also choose an auto reboot through the SCM.

Recovering from a System Failure

There are three kinds of system failures: power failures, hangs, and panics. The following sections describe how to recover from these failures.

Power Failure

Following a power failure, as soon as power returns to the system, it will reboot without operator intervention. See "Restoring File Systems after a Failure" later in this chapter. If your system has the uninterruptible power supply (UPS) subsystem, see "Managing the Uninterruptible Power Supply Subsystem" later in this chapter.

Hang

A hang occurs when an undetected condition causes system activity to halt, effectively freezing every process on the system. You regain control of the system by entering the hot-key sequence at the system console. To type the hot-key sequence hold down the Ctrl key while typing the following series of six bracket characters:

```
] [  ] [  ] [
```

Another way to express the same series is like this:

```
<Ctrl-]> <Ctrl-[> <Ctrl-]> <Ctrl-[> <Ctrl-]> <Ctrl-[>
```

This sequence induces a panic condition. See "Responding to a Panic" in the next section to recover from a panic.

If the hot-key sequence fails to induce a panic, use the hardware reset switch to interrupt the hang. Resetting the hardware restores control to the SCM, where you can reboot the system or, if you want the Data General Customer Support Center to investigate the cause of the problem, proceed to make a tape containing information required for diagnosis. See "Completing the Diagnostic Tape" later in the chapter.

If the reset switch fails to restore control to the SCM, turn off power to the computer and power the system back up. See "Restoring File Systems after a Failure."

Panic

A panic is a condition detected by the kernel that indicates a fatal malfunction or internal software inconsistency. Upon detecting a panic condition, the kernel halts all activity on the system and displays a message like the following at the system console:

```
DG/UX System Panic.  Panic code  57000072
```

If you wish to know the nature of the panic after you have restored the system, see the files in **/usr/release** that list panic codes. For example, to find the file containing information on panic code 57000072, you would execute the following commands:

```
# cd /usr/release ↵
# grep 57000072 *panic.codes ↵
```

You then read the indicated file using a command such as **view** or **more**, both of which offer commands for locating the desired text. See the **view**(1) or **more**(1) manual page.

Workstations that do not have error-correcting memory may occasionally experience an Unrecoverable Memory Error panic, code 1000037. This panic does not necessarily indicate that you have a bad memory module; it simply occurs occasionally on systems without error-correcting memory. Rectify the problem by turning the workstation off and waiting at least one minute before powering it back up and rebooting. If the panic occurs often (more than once every few months), you should call your Data General representative and have your system's memory modules tested and replaced if needed.

Responding to a Panic

The first thing to do when a panic occurs is to record the panic code number in the system log. Depending on how you have configured your system, the system may have proceeded beyond the panic code report. In any case, you then have several options:

- You may dump system memory to tape or disk so that the Data General Customer Support Center can help you diagnose the problem.
- You may halt the system, leaving it at the SCM.
- You may reboot the system.
- You may shut off power to the system (can occur without operator intervention on selected computers).

By default, the system responds to a panic by displaying the panic code and then prompting you to take a system memory dump. If you choose to take the system dump, the system issues several prompts for information before starting the dump. After the dump completes, or if you entered **no** to the dump prompt, the system restores control to the SCM. You may then reboot the system. If you wish, you can change any of the default system behavior described here using the **dg_sysctl**(1M) command.

The following sections discuss these options and tell how you can use the **dg_sysctl** command to configure your system to respond to a panic in a variety of ways, including recovering completely without operator intervention. If you do not wish to make a dump tape for diagnosis, see the following section.

Skipping the System Dump Procedure

If you do not wish to take a system dump for submission to the Data General Customer Support Center for diagnosis, you enter **no** at the system dump prompt that appears after a panic.

NOTE: We recommend that you take a system dump after every failure resulting from the DG/UX system software and, after copying the system dump and kernel image to tape, submit the tape to Data General for diagnosis. This information is very important to the Customer Support Center in their diagnosis and resolution of your problem.

To set up your system to skip the system dump without operator intervention after a panic, issue the **dg_sysctl** command with the **–d** option, as follows:

```
# dg_sysctl -d skip ↵
```

Thus configured, your system will respond to a panic by either halting or rebooting, according to how you set the automatic boot feature discussed later. If you configure your system to skip the system dump, you may skip the following section.

Taking a System Dump

To investigate the cause of a panic or hang, the Data General Customer Support Center requires a tape containing at least two files with the following contents:

File 0: the system memory contents, called the *system dump*.

File 1: the kernel executable, typically **/dgux**.

You may also append other files to the tape if you suspect that they may have contributed to the failure. Once you have restored the system, you complete the tape by dumping your kernel executable (and other relevant files, if any) to the tape. A later section, "Completing the Diagnostic Tape," tells how to make the tape for diagnostics.

Setting up an Automatic System Dump

A system dump is either **automatic** or **interactive**. An automatic dump occurs if you have used the **–d** option of **dg_sysctl** to set the the DG/UX automatic dump feature. This feature allows the system to initiate the dump sequence after a panic without operator intervention. The following command line enables the automatic dump feature:

```
# dg_sysctl  -d  auto ↵
```

Setting your system for automatic dump allows you to reduce the recovery time after a panic, but it also implies that you must make sure that the destination dump device is ready at all times to receive the system dump in the event of a panic. A later section describes the dump destination device in more detail.

Setting up an Interactive System Dump

An interactive dump is your system's default response to a panic. You can set this behavior explicitly by issuing the following command line:

```
# dg_sysctl  -d  ask ↵
```

When the system is configured to perform an interactive dump, the system responds to a panic by displaying the panic code and then the following prompt:

```
Do you want to take a system dump? [Y]
```

If you press Enter, the system prompts for the dump type and the dump destination device, described in the following sections.

Starting a Dump from the SCM

On systems where you had to respond to a hang by pressing the hardware reset switch, you may initiate an interactive dump by entering the following command at the SCM prompt:

```
SCM> START  1000 ↵
```

The system then begins prompting for the required information.

NOTE: The **START 1000** command does not produce a useful system dump if you have turned off power since the panic occurred or have already produced a system dump with **START 1000** since the panic occurred. In either of these cases, you cannot produce a useful tape for diagnosis, and you may proceed to reboot your system. See "Rebooting after a System Failure."

The dump procedure requires that you decide what type of dump to take and to which device to write the dump.

Selecting a Dump Type

You may dump either the entire memory of your system or just the kernel memory. A dump of just kernel memory is sufficient to diagnose a hang or panic unless your Data General representative tells you otherwise. The kernel memory dump, which is the default type, is faster and only requires around half the space of a complete dump. To change the default dump type to a complete dump, use the **dg_sysctl** command with the **-l** option:

```
# dg_sysctl  -l  all ↵
```

To restore the dump type to the default, substitute **kernel** for **all** in the command line above.

Selecting a Dump Destination Device

The dump destination is the tape device, virtual disk, or network interface to which you wish to write the dump. By default, the device is the value of the DUMP tunable parameter configured in your kernel. To change the value of this parameter, see Chapter 4. You can override the DUMP parameter setting with the **dg_sysctl** command's **-f** option, for example:

```
# dg_sysctl  -f  "st(ncsc(),5)" ↵
```

You may combine the **-d** and **-f** options to set the automatic dump feature and the dump destination device in a single command line, for example:

```
# dg_sysctl  -d  auto  -f  "st(cisc(),4)" ↵
```

OS Client Dump Destination Device

OS client systems should dump to their network interface, **inen()**. The OS server receives the dump over the network and writes it to a file created for the purpose by

the **sysadm** operation Client –> OS –> Add . By default, the Client –> OS –> Add operation creates an empty dump file for the client called **/srv/dump/***client_name*, which it lists in the client's **/etc/bootparams** entry on the server. The operation also exports the dump file for **root** access by the client, adding the appropriate entry to the server's **/etc/exports** file. The system administrator of the OS server should verify that the file system containing the dump file has sufficient free space. Use the **df**(1M) command to display the amount of free space in a file system. For more information on OS client setup, see Chapter 6.

Dumping to a Virtual Disk

A system with a local disk can dump to a local virtual disk instead of a tape. The advantage of dumping to disk is that it is faster than dumping to tape, resulting in decreased down-time. The disadvantage is that you must reserve for the purpose a virtual disk large enough to contain the system dump image. The dump image is equal to the size of the computer's physical memory plus 5 percent (if a complete dump), or around half that size (if a kernel dump).

NOTE: You can dump to a virtual disk only if it resides entirely on a local SCSI disk. You cannot dump to a virtual disk that comprises multiple partitions spanning multiple physical disks or to any virtual disk residing on an SMD or ESDI physical disk.

Create the virtual disk with the **sysadm** operation Device –> Disk –> Virtual –> Create. It is a good idea to give the virtual disk an appropriate name such as **sys_dump**. The dump process will write over any data, such as a file system, that resides on the virtual disk at the time of the dump. Therefore, you should not create a file system on the disk and attempt to use it for any purpose other than to contain the dump. By default, the system displays a prompt at the system console before writing to the dump virtual disk, allowing you to specify a different virtual disk or a tape drive if preferred. When the system panic procedure prompts you for the dump destination device, specify the physical and virtual disks using the following syntax:

vdm_dump(*physical_disk_name*,***virtual***_*disk_name*)

For example, use the **dg_sysctl** command and the **–f** option to set the dump destination device to virtual disk **sys_dump** on local physical disk **sd(cisc(),1)**:

```
# dg_sysctl -f "vdm_dump(sd(cisc(),1),sys_dump)" ↵
```

Rebooting after a System Failure

AV8500 and AV9500 systems have a "re–powerup on panic" feature, which determines whether or not the system will be booted automatically following a panic. Enabled by default, this feature overrides the **dg_sysctl** and SCM autoboot settings for the action to be taken after a system failure.

For all other systems, following a system failure, the processor is halted and control is restored to the SCM by default. From the SCM you may then reboot the system using the SCM **BOOT** command.

To minimize your system's downtime, you may set up your system to reboot without operator intervention. Set the automatic boot feature using two operations: the

dg_sysctl command **-r** option and the SCM's "Auto–reboot/Boot on error" flag, as follows:

```
# dg_sysctl -r auto

SCM> f

From the main menu, select "Change default boot paths" and
set the menu item "Auto-reboot/Boot on error" to "Enabled."
```

These settings configure your system to reboot whether or not a system dump is taken. By default, the automatic boot feature reboots the system using the most recently used **BOOT** command line, the one last used before the panic occurred. For example, if you last booted **sd(cisc(),0)root:/dgux.test**, the system will reboot using this boot path. To override this default behavior, use both the **dg_sysctl** command with the **-b** option and the SCM "Change default boot paths" menu to specify a different boot path. For example, both the following **dg_sysctl** command line and a matching boot path selected through the SCM "Change default boot paths" option, sets up the system to take a system dump and then boot from **sd(cisc(),2)root:/dgux**:

```
# dg_sysctl -d auto -f "st(cisc(),5)" -r auto -b
"sd(cisc(),2)root:/dgux"
```

You can set up your system to send you mail every time it reboots. This capability is particularly helpful because it reports reboots that occurred in your absence. To enable this feature, edit the **/etc/dgux.params** file. Set the **reboot_notify_START** parameter to true, and set the **reboot_notify_ARG** parameter to one or more local mail addresses. A notification message is sent to each user specified by the **reboot_notify_ARG** parameter each time the system boots.

Completing the Diagnostic Tape

Earlier in the recovery process, you may have decided to take a system dump so that you can make a tape for diagnosis by the Data General Customer Support Center. This section tells how to finish preparing the tape. The release notice, on-line in **/usr/release**, also tells how to prepare the tape.

The tape needs to be in the following format:

File 0: the system memory contents, called the *system dump*

File 1: the kernel executable, typically **/dgux**

The tape may also include other files if you suspect that they may have contributed to the failure.

NOTE: The diagnostic tape *must* include both the system dump and the kernel. The tape is useless without both of these images.

At this point in the recovery process, you need to copy the system dump to tape as file 0 from one of the following three places:

Tape If you wrote the system dump to tape after the failure, it is already on the tape as file 0.

A virtual disk

If you wrote the system dump to a virtual disk, transfer the system dump to a tape using the **lsd**(1M) command. For example, if your virtual disk is named **sys_dump**, use the following command line to dump it to the tape at **/dev/rmt/0**:

```
# lsd -t /dev/rdsk/sys_dump /dev/rmt/0n ↵
```

The command line above does not rewind the tape after writing to it.

A file on the OS server (OS clients only)

By default, the Client –> OS –> Add operation creates **/srv/dump/***client_name* as the destination for a system dump by client *client_name*. If you specified a different dump destination when adding the OS client, look there instead. An OS client dumps to the file named in its **/etc/bootparams** entry on the OS server. To transfer the dump file to a tape, use the **cpio**(1) command with the **–oBcv** options. For example, to write a system dump file from OS client **ralph** to the tape at **/dev/rmt/0**, execute the following commands:

```
# cd /srv/dump ↵
# echo ralph | cpio -oBcv > /dev/rmt/0 ↵
```

You should dump the file from the directory where it is located using a relative path name as shown above. Do not dump the file by passing an absolute pathname to **cpio**.

Dumping the Kernel Executable

With the system dump written to tape, you are ready to dump the kernel executable. Be careful not to overwrite the system dump when you dump the kernel; if you overwrite the system dump , the Data General Customer Support Center cannot diagnose your problem. To make sure the tape is positioned at the end of the system dump (file 0), use the **mt**(1) command to position the tape. For example, to position the tape in tape drive **/dev/rmt/0** at the end of file 0, issue the following commands:

```
# mt -f /dev/rmt/0 rewind ↵
# mt -f /dev/rmt/0n fsf 1 ↵
```

If there is not room on the tape for the kernel executable, you may write it to a second tape instead.

By default, the kernel executable is **/dgux**. If the failure occurred with a different kernel, however, you should dump it instead of **/dgux**. If dumping an OS client's kernel from the OS server, be careful to dump the correct file. The file that appears as **/dgux** in an OS client's file system appears as **/srv/release/***release_name***/root/***client_name***/dgux** on the OS server. For example, if you want to dump the kernel that OS client **ralph**, who uses the primary release, refers to as **/dgux**, you should dump **/srv/release/PRIMARY/root/ralph/dgux**.

To dump the kernel, use the **cpio**(1) command with the **–oBcv** options. For example, the following command line dumps **/dgux** to the tape at **/dev/rmt/0**, rewinding the tape when done:

```
# cd / ↵
# echo dgux | cpio -oBcv > /dev/rmt/0 ↵
```

You should dump the file from the directory where it is located using a relative pathname as shown above. Do not dump the file by passing an absolute pathname to **cpio**.

Dumping other files to tape
If you suspect that other programs or files may have contributed to your system's failure, you may include them on the tape as well. You may dump these files as tape files 2, 3, and so on. As when dumping the kernel executable, dump them using **cpio –oBcv** using relative path names.

Labeling the Tape
When you have finished making the tape, label it specifying the name of your company, the cause of the failure, the date, and the contents of the tape. Your label might look like this:

```
BLUE DAEMON SYSTEMS, INC., Durham, NC
Panic code:  3400002
Date:  April 6, 1993
File 0:  system memory dump
File 1:  kernel executable
File 2:  miscellaneous files
Density:  QIC-525 tape at high density
cpio format:  cpio -oBcv
```

Halting the System

To halt your system immediately after a panic without taking a system dump, enter **no** at the prompt to take a system dump. If you take the default (**yes**) instead, the system by default takes the system dump and halts. You may then boot with the SCM **BOOT** command.

Using the **dg_sysctl** command's **–r** option, you can set up your system to halt after a panic or system dump without operator intervention. To set up your system to halt after a panic without taking a system dump, issue the following command:

```
# dg_sysctl -d skip -r halt ↵
```

To set up your system to take the system dump without operator intervention and then halt, issue the following command:

```
# dg_sysctl -d auto -f "st(cisc(),5)" -r halt ↵
```

Shutting off Power to the System

Some Data General AViiON computers support a feature allowing you to shut off power to the system using a software command. Using this feature, you can set the

system to shut off power after a normal shutdown. You cannot set the system to shut off power after a panic.

On systems that support this feature, use the **-p** option of the **dg_sysctl** command to set up the system to shut off power after normal shutdown. For example:

```
# dg_sysctl -p auto ↵
```

The **dg_sysctl** command ignores the **-p** option if your system is set for automatic boot (as with **-r auto**). On systems that support the power-off feature, the default is **auto**; on systems that do not support this feature, the default is **skip**.

Restoring File Systems after a Failure

A failure on a DG/UX system may not damage files on your system. Damage does occasionally, occur, however, resulting in file system metadata becoming inconsistent or data being lost. The first time you boot your system after a failure, the DG/UX system performs several operations intended to seek out and, where possible, repair damage to files and file systems.

By default, the system invokes **fsck** to check the / file system upon rebooting, both after a system failure and during normal operation. If you do not want this initial check to occur, or if you want to change the kind of check that **fsck** performs, change the RUNFSCK and FSCKFLAGS tunable parameters in the system file you use to build your kernel. See the section on setup and initialization configuration variables in Chapter 4 for more information.

If the system failure damaged files necessary for bringing up the system, **fsck** may fail. If this happens, see "Repairing Damaged DG/UX System Files."

Once the system has booted successfully, it will proceed to check local file systems according to their **fsck** pass numbers, which you may review with the **sysadm** operation `File System -> Local Filesys -> List`. You can speed up this process by mounting your file systems for fast recovery. For more information on fast recovery file systems and **fsck**, see the **fsck** section in Chapter 9.

When file system checking is complete, you should check **/etc/log/fsck.log** and all local file systems' **lost+found** directories to see if you need to restore any lost or damaged files from backup. To restore files from backup, see Chapter 9.

The **fsck** utility has no way of verifying the contents of files on your system. If you use a database product for example or some other software that can check the files it uses, you may want to invoke them for this purpose. The **fsck** utility can verify only the structure of the DG/UX file system.

Repairing Damaged DG/UX System Files

This section describes how to recover after **fsck**(1M) fails to repair the / or **/usr** file system at boot. As a result, the system will not boot. One common error that makes it impossible to boot a system is when the **/etc/inittab** file is damaged. When this happens, you frequently see a message like this when **init** starts running:

```
... SAD autopush configuration failed ...
```

If a system or disk failure damages **inittab** or other DG/UX system files (those in the **/usr** or / file systems), you need to repair the file systems and restore the damaged files from backups. If you cannot repair the file system, you need to reload the system software from either the release media or a bootable copy of the file systems made with the **systemtape** utility.

If you are running DG/UX 5.4R3.00 or later, you can use the CD–ROM release disk to repair your system. The CD–ROM can be booted into repair mode, under which a complete working environment is available from file systems on the CD. To boot the CD into repair mode, specify the **–R** option flag in the SCM boot command line:

```
SCM> b sd(cisc(),3) -R
```

After you have the repair mode environment running, repair the damaged / and **/usr** file systems' virtual disks with the **fsck** command and then mount them on different mount point directories. You can then copy files from the CD or from system backups to repair any damage.

If you do not have the CD-ROM release, you can use stand–alone **sysadm** to repair the file systems. Boot stand-alone **sysadm**, either from **/usr** or the release tape.

If you want to boot from **/usr**, use a command line like the following, where you specify the physical disk on which the **/usr** file system resides:

```
SCM> b sd(cisc(),0)usr:/stand/sysadm ↵
```

An attempt to boot stand-alone **sysadm** will fail if the **/usr** file system is corrupt or if the **/usr** file system is built on a virtual disk consisting of multiple partitions spanning multiple physical disks. If stand-alone **sysadm** fails or has been deleted from **/usr**, then boot the release tape for the revision your system is running using a command line like the following:

```
SCM> b st(cisc(),4) ↵
```

If stand-alone **sysadm** boots successfully, go to the File System Management Menu and use the Check a File System operation to check the / and **/usr** file systems. When prompted for **fsck** flags, specify **–xlp**.

If you still cannot boot the system, use the Check a File System operation again, but specify the **–y** option for **fsck**. The **–y** option repairs all non-fatal flaws in the file system, even if the repair results in lost files or data.

If **fsck –y** fails, use the stand–alone **sysadm** installation procedure to get the damaged / and **/usr** files systems mounted. Execute Install Software –> Prepare Virtual disks. Answer the system queries as follows:

```
2.  Prepare required virtual disks

Run this step now?  [yes] ↵
Required File System Mount Points:
```

```
File System    Virtual  Current   Action     Blocks  Physical
Mount Point    Disk     Blocks    Required   To Add  Disk
-------------  -------  -------   --------   ------   --------
-none-         swap      50000      None        -    sd(insc(0),0,0)
/              root      40000      None        -    sd(insc(0),0,0)
/usr           usr      240000      None        -    sd(insc(0),0,0)

Modify this information? [no] ↵
```

At this point, escape to the shell by entering **!** at the **sysadm** prompt. The damaged / file system is mounted as **/mnt/root**. the damaged **/usr** file system is mounted as **/mnt/usr**. Set your path as follows to gain access to the normal commands:

```
# PATH=/mnt/root/sbin:/mnt/usr/bin:/mnt/root:sbin↵
# export PATH ↵
```

Note that some commands may not work if they use shared libraries, since the real **/usr** file system that contains them is mounted in a different place than normal. You can now change directory to **/mnt/root** and **/mnt/usr** to examine and repair the damaged file systems. If you are repairing a damaged **inittab** file, there may be an undamaged backup of it in **/etc/inittab.backup**. The original prototype file is in **/etc/inittab.proto**.

When you finish fixing the file systems, exit the shell by entering **exit**. This puts you back in **sysadm**. Exit **sysadm** by entering **q**. That shuts the system down. You can now reboot your repaired / and **/usr** to bring the system back up.

If **fsck –y** succeeds, try to boot the system. If the boot fails, you must boot the release tape and reload the DG/UX system files by going to the Install System Software Menu and performing the Load software operation. The Load software operation will load system files as necessary. When it has finished, you may replace any files you wish by restoring them from backup.

If all of these options fail, you must re-install the DG/UX system completely. Do this by going to the Virtual Disk Management Menu and using the Delete a Virtual Disk operation to delete the virtual disks containing the / and **/usr** file systems (**root** and **usr** virtual disks). Then go to the Install System Software Menu and select option 7, All steps. You will also need to re-install any packages that had been installed before the failure.

As an alternative to booting from the release tape, you can use a bootable copy of your / and **/usr** file systems made with the **systemtape** utility. In addition to normally scheduled backups, you can use **systemtape** to create a bootable tape of these file systems at regular intervals. You can then use that tape if you need to restore your system. See the **systemtape**(1M) manual page for more information.

Logging System Errors and Messages

Various system facilities produce messages during normal operation and when they encounter errors and other unexpected conditions. **Sysadm** automates a procedure for using the system error log. See Chapter 16 for more information.

Managing the Uninterruptible Power Supply Subsystem

This section applies only if you have installed the uninterruptible power supply (UPS) subsystem on your computer. The UPS subsystem monitors the power supply to which your system is connected. If the power supply fails, the UPS subsystem provides a limited additional power supply. Depending on how you configure the UPS daemon running on the host, it may perform shutdown and/or reboot of the host depending on the state of the power supply and the UPS backup battery.

Specifically, the UPS subsystem functions as follows. Once set up, the UPS daemon starts at boot and polls the backup battery unit at regular intervals (30 seconds by default). The daemon polls for two pieces of information: the status of the line power supply, and the status of the backup battery. This is how the system functions under normal circumstances.

When line power fails, the UPS backup battery supplies power to the host. The next time the UPS daemon polls the backup battery unit, it detects that line power has failed and does the following on the host:

1. The UPS daemon logs a system error using the **syslog** error logging facility. By default, **syslog** responds to the UPS message by alerting all users of the power failure. You can change this behavior by editing **/etc/syslog.conf** and **/usr/sbin/ups.action**. See **syslog.conf**(5) for more information.

   The UPS daemon invokes **/usr/sbin/ups.action** under either of these conditions:

   - Line power has terminated.
   - Line power returns before the battery dies.
   - Battery power is low.
   - Line power returns when the system is on battery power.

   This script uses the **wall**(1M) command to send commands to all users. If you do not want to notify all users, you may edit the script to redirect messages to the system console or elsewhere.

2. The UPS daemon begins counting down a time-out sequence before shutting down the system. You may define the length of the time-out sequence using a **sysadm** operation to be discussed later.

If line power returns before the UPS daemon begins the shutdown, the daemon aborts the time-out and the backup battery unit restores the normal line power supply to the computer. If line power does not return before the time-out ends, however, the UPS daemon shuts the system down to single-user mode.

The system remains in single-user mode either until line power returns or until the battery fails and the system powers down. If line power returns before the battery fails, the UPS daemon takes the system back up to the default run level. If line power returns after the battery has failed, the system reboots.

The operations for managing the UPS subsystem are in the **sysadm** menu `Device -> UPS` and `Availability -> UPS`. These operations call the **admups**(1M)

command to perform the requested operation. The **admups** command may also offer additional functions.

Setting up the UPS Subsystem

To use the UPS subsystem, you need to dedicate a terminal line for connection to the UPS backup battery unit. The UPS daemon running on the host uses the terminal line to communicate with the backup battery unit. Connect the UPS backup battery subsystem to the port using the Data General cable supplied specifically for this purpose.

The serial port must include modem control. The serial port must not have a port service currently assigned to it. It is not sufficient simply to disable the port service on the port; you must delete the port service with the **sysadm** operation `Device -> Port -> Port Service -> Delete`. See Chapter 10 for more information on port services.

NOTE: If you start a port service on the port selected for the UPS system, unpredictable results could occur, possibly requiring sudden shutdown of the system.

To set up the UPS subsystem, connect the UPS power cables according to the UPS hardware documentation. Connect the communication cable to the UPS backup battery unit and the selected serial port on the computer host. It is important that you complete installation of the UPS subsystem before starting the UPS daemon.

NOTE: If you start the UPS daemon before you have installed the UPS backup battery unit, unpredictable results could occur.

Execute the **sysadm** operation `Device -> UPS -> Start` or `Availability -> UPS -> Start`. This operation starts the UPS daemon and sets up your system to start the daemon at boot. The operation prompts for the following information:

`Polling Interval` This parameter determines how often, in seconds, the UPS daemon polls the backup battery unit for the status of line power and backup battery viability. In effect, this value determines the maximum number of seconds that may pass between a line power failure and the beginning of the time-out countdown initiated by the UPS daemon. By default, the UPS daemon polls the backup battery unit every **30** seconds.

`Timeout` This parameter determines how many seconds the UPS daemon waits before commencing shutdown after detecting that line power has failed. If the parameter is **0**, the UPS daemon waits indefinitely, delaying the shutdown until it detects that the backup battery unit is low on power. The backup battery unit is considered low on power when it has approximately 2 minutes of power left. By default, the parameter is set to **0**.

`Serial Port` This parameter is the pathname of the serial port that the UPS daemon will use to communicate with the backup battery unit. The port must include modem control. There should be no port service, whether enabled or disabled, associated with this port. See Chapter 10 for more information on port services. Setting this parameter to **none** effectively disables the UPS daemon and makes the port available for use with a terminal, modem, or printer.

Stopping the UPS Subsystem

To stop the UPS daemon and change your system setup so that the UPS daemon no longer starts at system boot, select the operation `Device -> UPS -> Stop`.

To restore use of the dedicated serial port as a normal terminal line, follow these steps:

1. Use the **sysadm** operation `Device -> UPS -> Parameters -> Set` to reset the daemon's serial port parameter to another port or to the value **none**. Setting the port to **none** stops the daemon and disables it from starting at boot.
2. Remove the specialized UPS cable connected to the port. Refer to the UPS hardware documentation before disconnecting the UPS backup battery unit.
3. Use the sysadm operation `Device -> Port -> Terminal -> Add` to start the **login** service for the serial port.

Setting and Displaying UPS Parameters

The **sysadm** menu `Device -> UPS -> Parameters` contains the Set operation for setting the values of the UPS parameters. You do not need to stop the UPS daemon to change its operating parameters. Use the List operation to display the values of the parameters. For a discussion of the parameters, see "Setting up the UPS Subsystem."

Displaying UPS Status

Use the **sysadm** operation `Device -> UPS -> List` to display the status of the UPS subsystem. The display indicates:

- Whether the UPS daemon has detected a line power failure.
- Whether the backup battery unit has reported that it is approaching failure.
- Whether the UPS daemon has initiated the shutdown sequence.

Optionally, you can also use the operation to review the history of UPS events as appearing in the system log maintained by the **syslog** system error logger.

Expert Information

You do not have to read this expert section to operate the DG/UX system. Information here is optional and is provided to enhance your understanding of how

run levels work. For basic information about run levels, see "DG/UX System Run Levels" earlier in the chapter.

The Fundamentals

The **inittab** file contains entries specifying which processes will be invoked at which run level.

The **init** program reads the entries in **inittab**, and when they match the specified run level, **init** passes them to a shell for execution.

The **rc.init** script, when called with an argument S through 6, executes the scripts in the given **rc***N***.d** directory. The processes are invoked according to **K** (kill) and **S** (start) switches.

The scripts in the **rc***N***.d** directories are **rc** scripts. Commonly called "run command" scripts, they start and stop system services required by run levels S through 6. Output from **rc** scripts goes to **/etc/log/init.log**.

The **rc***N***.d** directories are used to organize and order the set of run command scripts associated with a particular run level. To avoid duplicate scripts and the problem of maintaining consistency among duplicate scripts, the entries in an **rc***N***.d** directory are links to a specific run command script in **init.d**.

The **init.d** directory contains all of the run command scripts. Some are started at many run levels; some are started at one level and stopped at all other levels; some are started and never stopped until reboot time.

The Sequence

Let's follow the sequence that occurs when you invoke **init** to set a run level. Assume your system has been booted and is going to be changed from single-user mode, run level S, to multiuser mode, run level 3. We'll track one of the processes invoked, **syslog**.

1. Invoke the **init** program with the argument 3:

   ```
   # init 3 ↵
   ```

2. The **init** program scans the **inittab** file for all entries containing the run level number 3 in the *run level* field. Then, **init** invokes the `rc.init 3` instruction which is in the *process* field.
3. The **/sbin/rc.init** program uses the run level number 3 as a pointer to directory **/etc/rc3.d**, which contains links to the scripts in **/usr/sbin/init.d**. A script called **rc.syslogd** starts the **syslogd** program.
4. The **rc.init** program then executes all scripts for run level 3; among these is **syslogd**.

The /etc/inittab File

The **init** program relies on the information in the **/etc/inittab** file, whose entries have this format:

id:level:action:process

where the fields are as follows:

id one, two, or three characters that uniquely identify an entry.

level a character (**s**, **0** through **6**, **a**, **b**, or **c**) that determines at what run level the specified action is to take place. If the level field is empty, the action occurs at all run levels.

action
one of the following:

boot Process the entry the first time **init** leaves single-user mode. Do not wait for the process to terminate.

bootwait Process the entry the first time **init** goes from single-user mode to multi-user mode after the system is booted. If **initdefault** is set to 2, the entry is processed at boot time. **init** starts the process, waits for its termination and, when it dies, does not restart the process.

initdefault
When **init** starts, it will enter the specified level. The process field for this action is not used.

off At the specified level, kill the process or ignore it.

once Run the specified process once and don't start it again if it finishes.

ondemand Synonymous with **respawn**, but used only when the level is **a**, **b**, or **c**.

powerfail Execute the process in this entry only when **init** receives a power fail signal (SIGPWR). See **signal**(2).

powerwait Execute the process in this entry only when **init** receives a power fail signal and wait until it terminates before continuing any processing of **inittab**.

respawn If the process does not exist, start it, wait for it to finish, and then start another.

sysinit Process the entry before **init** attempts to access the system console. Wait for the process to terminate before continuing.

wait When going to the specified level, start the specified process and wait until it's finished.

process
any executable program, including shell procedures.

You can add a comment to the end of a line by preceding the comment with a pound sign (#). The **init** program ignores everything appearing after a pound sign on a line.

Now let's look at the prototype **inittab** file and see how the structure makes sense to the DG/UX system:

```
#
def:s:initdefault:
ttc::sysinit:/sbin/autocon                </dev/console >/dev/console 2>&1
fsc::bootwait:/sbin/chk.fsck              </dev/console >/dev/console 2>&1
dat::bootwait:/usr/sbin/init.d/chk.date   </dev/console >/dev/console 2>&1
set::bootwait:/usr/sbin/init.d/chk.system </dev/console >/dev/console 2>&1
tty::bootwait:/usr/sbin/init.d/chk.strtty </dev/console >/dev/console 2>&1
dev::bootwait:/usr/sbin/init.d/chk.devlink</dev/console >/dev/console 2>&1
#
rc0:0:wait:/sbin/rc.init 0 >/dev/console 2>&1
rci:i:wait:/sbin/rc.init i >/dev/console 2>&1
rc1:1:wait:/sbin/rc.init 1 >/dev/console 2>&1
rc2:2:wait:/sbin/rc.init 2 >/dev/console 2>&1
rc3:3:wait:/sbin/rc.init 3 >/dev/console 2>&1
rc4:4:wait:/sbin/rc.init 4 >/dev/console 2>&1
rc5:5:wait:/sbin/rc.init 5 >/dev/console 2>&1
rc6:6:wait:/sbin/rc.init 6 >/dev/console 2>&1
#
# ttymon is more secure than su since su is always on console
con::respawn:/usr/lib/saf/ttymon -g -p "Console Login: " -d /dev/console \
        -l console
sec::off:#/sbin/su - 1 </dev/console  >/dev/console 2>&1
#
saf:234:respawn:/usr/lib/saf/sac -t 45  #Service Access Facility

# When init receives a SIGPWR, it kicks itself with 'init S':
ups::powerfail:/sbin/init S < /dev/console > /dev/console 2>&1
```

Figure 3–1 The Prototype /etc/inittab File

The line beginning with `con` was broken for readability.

The first line in the file sets `s`, single-user mode, as the default initialization run level.

The next line makes the system console usable by pushing the required STREAMS modules (for line discipline and so on) onto the stack that controls the system console.

The next five lines start up five check scripts: **chk.fsck, chk.date, chk.system, chk.strtty, and chk.devlink**. These scripts are executed at boot time according to the **bootwait** action of **inittab**(4).

The next eight lines are instructions for setting run level **i** (installation) and run levels 0 through 6. For instance, at run level 3, the **init** program invokes the **rc** scripts in **/etc/init.d** via the links in **/etc/rc3.d**. These scripts perform the functions necessary to start system services for run level 3, and to stop services not associated with run level 3. Standard output and standard error are directed to **/dev/console** for all run levels.

After the run level lines, the line `con` identifies the operator's console (**/dev/console**) to the system. The line after it, `sec`, is an alternate (less secure) service to run on the system console. Note that this line is turned off. The next line, `saf`, starts the Service Access Facility (SAF), at run levels 2, 3, and 4 to provide

terminal services to users. The last line, `ups`, shuts down the system if a powerfail condition occurs. This service exists only on systems with the uninterruptible power supply (UPS) subsystem, described earlier in the chapter.

RC Scripts and Check Scripts

When **rc.init** invokes a run level, the characteristics of that run level are produced by scripts in **/usr/sbin/init.d**. There are two types of scripts:

chk.* These scripts are usually run once, at boot time. An example is **chk.fsck** which runs the **fsck** program on file systems.

rc.* These scripts are invoked with either a start or stop argument. An example is **rc.tcload**, which starts or stops the asynchronous terminal I/O controllers.

Init.d Links

Typically, the above scripts exist in **/usr/sbin/init.d**. The **rc** scripts are invoked via links in an **/etc/rc*N*.d** directory. Remember, there are nine **/etc/rc*N*.d** directories: **/etc/rcS.d**, **/etc/rc0.d**, **/etc/rci.d**, **/etc/rc1.d**, **/etc/rc2.d**, **/etc/rc3.d**, **/etc/rc4.d**, **/etc/rc5.d**, and **/etc/rc6.d**. The *names* of the links are labeled as follows:

`Snnn.name`

or

`Knnn.name`

The entries have three parts:

`S` or `K`	Defines whether the process should be started (`S`) or killed (`K`) upon entering the new run level.
nnn	A number from 000 to 999 indicating the order in which the files will be started (S111, S112, S113, and so on) or stopped (K231, K232, K233, and so on).
name	The script name in **/usr/sbin/init.d**.

All process scripts are specified to be either killed or started when you change run levels. The **rc.init** program executes all **K** scripts first; they are executed from highest ID number to lowest ID number. When all **K** scripts have executed, **S** scripts begin executing from lowest ID number to highest ID number. All scripts in **init.d** have links in all **/etc/rc*N*.d** directories; the **K** or **S** prefix determines what is on and what is off.

For example, the run level 3 link name for **rc.localfs** is **S114.localfs**. This link is in **/etc/rc3.d**.

Let's look at the rest of **/etc/rc3.d**. Type:

```
# cd /etc/rc3.d ↵
# ls ↵
```

```
K237.ypserv    S116.sync      S212.llc       S239.nfslockd   S315.nfsserv
S015.ups       S117.lan       S232.tcpipport S251.account    S334.tcpipserv
S112.tcload    S119.setup     S235.syslogd   S252.cron       S353.nfsfs
S113.update    S130.daemon    S236.dgserv    S253.lpsched    S358.failover
S114.localfs   S210.usrproc   S237.ypserv    S254.preserve
```

The complete layout of how all **rc** scripts are started and killed is in **/etc/dgux.rclinktab.proto**. Figure 3–2 is a portion of that file:

```
#     run level        id   S   0   1   2   3   4   5   6   i

    rc.ups             015  S   S   S   S   S   S   K   K   -

##  rc.usrfs           111  K   K   S   S   S   S   K   K   -
    rc.tcload          112  K   K   S   S   S   S   K   K   -
    rc.update          113  K   K   S   S   S   S   K   K   S
    rc.localfs         114  K   K   S   S   S   S   K   K   S
    rc.sync            116  K   K   K   S   S   S   K   K   -
    rc.lan             117  K   K   K   S   S   S   K   K   -
    rc.setup           119  K   K   S   S   S   S   K   K   -
    rc.daemon          130  K   K   S   S   S   S   K   K   -
    rc.install         131  -   -   -   -   -   -   -   -   S
##  special systems    150

    rc.usrproc         210  K   K   K   S   S   S   K   K   -
##  rc.mcli            211  K   K   K   S   S   S   K   K   -
##  rc.eventd          2111 K   K   K   S   S   S   K   K   -
    rc.llc             212  K   K   K   S   S   S   K   K   -
##  rc.omtran          214  K   K   K   S   S   S   K   K   -
##  rc.netbeui         216  K   K   K   S   S   S   K   K   -
##  rc.x25port         217  K   K   K   S   S   S   K   K   -
##  rc.tsp             230  K   K   K   S   S   S   K   K   -
##  rc.tcpipport       232  K   K   K   S   S   S   K   K   -
    rc.syslogd         235  K   K   K   S   S   S   K   K   -
    rc.dgserv          236  K   K   K   S   S   S   K   K   -
##  rc.ypserv          237  -   -   -   S   S   S   -   -   -
##  rc.nfslockd        239  K   K   K   S   S   S   K   K   -
##  special systems    250
    rc.account         251  K   K   K   S   S   S   K   K   -
    rc.cron            252  K   K   K   S   S   S   K   K   -
    rc.lpsched         253  K   K   K   S   S   S   K   K   -
    rc.preserve        254  K   K   -   S   S   S   K   K   -
##  rc.nbvt            255  K   K   K   S   S   S   K   K   -
##  rc.nw_tran         270  K   K   K   S   S   S   K   K   -
##  rc.snalan          271  K   K   K   S   S   S   K   K   -
##  rc.sdlc            272  K   K   K   S   S   S   K   K   -
##  rc.icobol          293  K   K   K   S   S   S   K   K   -

##  rc.omserv          314  K   K   K   K   S   S   -   -   -
##  rc.nfsserv         315  -   -   -   K   S   S   -   -   -
##  rc.x25serv         317  -   -   -   K   S   S   -   -   -
##  rc.osit            318  K   K   K   K   S   S   K   K   -
##  rc.alp             319  K   K   K   K   S   S   K   K   -
##  rc.icl             330  K   K   K   K   S   S   K   K   -
```

```
##  rc.x400             331  K   K   K   K   S   S   K   K   -
##  rc.lanman           332  K   K   K   K   S   S   K   K   -
##  rc.tcpipserv        334  -   -   -   K   S   S   -   -   -
##  rc.smtpgw           336  K   K   K   K   S   S   K   K   -
##  special systems     350
##  rc.nfsfs            353  K   K   K   K   S   S   K   K   -
##  rc.upsd             355  K   K   K   K   S   S   K   K   -
##  rc.upsd.client      356  K   K   K   K   S   S   K   K   -
    rc.failover         358  K   K   K   K   S   S   K   K   -
##  rc.nw_serv          370  K   K   K   K   S   S   K   K   -
##  rc.sna              371
##  rc.sna              372
##  rc.sna              373
##  rc.x500             375
##  rc.cmips_agent      378
##  rc.oseyenode        379  K   K   K   K   S   S   K   K   -
##  rc.x11              391
##  rc.dni              395  K   K   K   K   S   S   K   K   -
##  rc.dims             396  K   K   K   K   S   S   K   K   -

    rc.halt             511  -   S   -   -   -   -   S   -   -

    rc.reboot           611  -   -   -   -   -   -   -   S   -
```

Figure 3–2 RC Scripts: the Kill and Start Mechanism

You can think of all **rc** scripts as being either on (S) or off (K). Since TCP/IP and ONC/NFS are optional products, we show their links commented out above. The **/etc/dgux.rclinktab.proto** file contains comments explaining this table.

Changing the Behavior of RC Scripts

The behavior of all **rc** scripts is governed by data and arguments set in a parameters file. There are several such parameters files:

- **/etc/dgux.params** (shipped with the DG/UX system)
- **/etc/tcpip.params** (shipped with TCP/IP)
- **/etc/nfs.params** (shipped with ONC/NFS)

For example, the **rc.nfsserv** script starts and stops the **nfsd** daemon. The more network interaction you have, the more copies of the daemon you would want. The parameter you would change in **nfs.params** is **nfsd_ARG**. To run twelve copies of the daemon, for instance, set the parameter to:

```
nfsd_ARG="12"
```

To add your own **rc** scripts, see *Porting and Developing Applications on the DG/UX™ System*.

End of Chapter

Chapter 4
Managing the System Configuration

The first part of this chapter lists the administrative logins for the DG/UX system and covers the following procedures:

- Recovering from a forgotten **root** or **sysadm** password
- Monitoring system activity
- Managing the accounting system
- Monitoring system process activity
- Managing security databases
- Building and executing a kernel
- Setting system parameters
- Setting the language environment
- Setting system date and time

The second part of this chapter offers suggestions for performance management, lists and definitions of the tunable parameters for the DG/UX system, and explains how to get information about the user license and upgrade user licensing.

The DG/UX Administrative Logins

On the DG/UX system, a file's owner controls access to the file. The superuser logins, **root** and **sysadm**, which exist on all systems, can override any permission settings and can execute, open, read, delete, or change any file in the system. If you know the **root** or **sysadm** password, you can become the superuser by executing the **su**(1) command with the - (hyphen) option followed by the desired login name. We recommend that you use **sysadm** rather than **root**, like this:

```
% su - sysadm ↵
```

When you log in as **sysadm**, your home directory is the **/admin** directory instead of /. Using this login keeps your personal administrative files out of /. As an alternative to using **su** to become **sysadm**, you can simply log in as **sysadm**.

The DG/UX system can convey system file ownership among 12 login names. Five of these login names function normally, that is, you can become these with **su**. The

other seven are for system use only; that is, you never actually log in with them. If you look at **/etc/passwd**, you'll see that the system logins have an asterisk (*) in the password field, meaning that no one can log in with these.

Generally, you should perform administrative tasks as **sysadm**. In some cases, however, you may need to use **su** to change to another system login name for some operations. For example, you should become **nuucp** before changing the **cron** jobs for UUCP. You do not want to become **uucp**, however, because that login is reserved for use by the UUCP facility's **uucico** daemon. For more information on UUCP, including a discussion of the difference between the **uucp** and **nuucp** logins, see *Using Modems and UUCP on the DG/UX™ System*. Table 4–1 shows the administrative and system logins.

Table 4–1 Default DG/UX Logins

Login	How It Is Used
root	This login has no restrictions. It overrides all process and file permissions. The **sysadm** login has the same unlimited access privileges as **root**.
xdm	Your system has this login only if you have installed the X11 package. Logging in as xdm executes telxdm(1X), the xdm control utility. See the telxdm(1X) manual page for more information. xdm is the X Window System session manager for systems with graphics capabilities.
sysadm	Same as **root**, except the login directory is **/admin**. You should use this login for performing administrative tasks.
daemon *	This is the login of the system daemon, which controls background processing.
bin *	This login owns the files in **/usr/bin**.
sys *	This login owns the files in **/usr/src**.
adm	This login owns the files in **/var/adm**.
uucp *	This login owns the object and spooled data files in **/usr/lib/uucp** and **/etc/uucp**. To make UUCP connections, systems log in to other systems with the UUCP login and initiate file transfers via **/usr/lib/uucp/uucico**.
nuucp	The system administrator can log in to the system as **nuucp** and perform general administrative tasks. Some UUCP facilities may send messages to **nuucp**'s mail file (**/var/mail/nuucp** by default).
lp *	This login owns the object and spooled data files in **/var/spool/lp**, **/etc/lp**, and **/usr/lib/lp**.

Continued

Table 4–1 Default DG/UX Logins

Login	How It Is Used
mail *	This is the login of the electronic mail facilities. Some system facilities may send messages to **mail**'s mail file (**/var/mail/mail** by default).
sync *	Logging in as **sync** causes the system to execute the **sync**(1M) command before returning you to the login prompt. The **sync** command no longer performs any function; therefore, logging in as **sync** has no effect at all. The **sync** command and the **sync** login exist only for compatibility with previous revisions of the system.

In Table 4–1, only those entries without asterisks may be used as actual logins; the others are for system use only. Additional software packages may add other logins to the **/etc/passwd** file.

Recovering Forgotten Superuser Password

If you forget both of the superuser (**root** and **sysadm**) passwords, follow the steps in this section to set a new one. You may be in either of two situations when you realize that you have forgotten the superuser passwords:

- You are logged on as **root** or **sysadm**, and you have the # prompt. Simply run the **passwd**(1) command and set a new **root** and/or **sysadm** password.
- You are not currently logged on as **root** or **sysadm**. For instance, you may be logged on as yourself and have the normal shell prompt. Because there is no way to access the **root** or **sysadm** logins, you will have to bring the system down in what is called an "unclean" halt.

If the latter condition is the case, go to the system console and do the following:

1. Use **wall**(1M) to warn users that the system is about to go down. Wait until users have logged off or have at least terminated any processes (such as editors) that write to disk files.
2. Wait 60 seconds before resetting your system. You reset your system either by pressing the reset button or by entering the hot-key sequence at the system console. The hot-key sequence consists of three pairs of square brackets ( ] [ ] [ ] [ ) pressed while you hold down the control key:

   ```
   <Ctrl-]>  <Ctrl-[> <Ctrl-]>  <Ctrl-[>  <Ctrl-]>  <Ctrl-[>
   ```

   Resetting the system takes you to the SCM prompt.

3. Reboot your system to bring it up to single-user mode. If you have configured your system to come up to a run level other than single-user mode, you need to specify the **-S** option on the SCM **boot** command line. For example,

   ```
   SCM> b sd(cisc(),0)root:/dgux -S ↵
   ```

4. If the attempt to boot fails because the root file system is corrupt, boot stand-alone **sysadm**, either from disk or from the release tape, and check the root file system. If you you need to repair damaged DG/UX system files, see Chapter 3.
5. Once you are in single-user mode, assign new **root** and **sysadm** passwords with the **passwd**(1) command.

Monitoring System Activity

The DG/UX system has a system activity reporting mechanism, which is described in manual pages **sar**(1) and **nsar**(1), that you can use to review statistics on CPU performance, disk and terminal I/O, memory usage, process communication and execution, and other activity. When the system appears to be functioning erratically, or simply as a matter of course, you may want to review the system activity data.

nsar (the new **sar** utility) provides additional system performance information such as paging and memory usage statistics, and information on virtual disks. **nsar** supports new options, one of which provides over one hundred new statistics.

The **sysadm** operation `System -> System Activity` menu provides operations for starting and stopping **sar** or **nsar**, deleting old data collections, and reviewing reports.

The **sar** and **nsar** utilities accumulate the data and produces reports. Depending on the one you use, it does not run continuously. Instead, you must start it, allowing it run, before you get data to review. When **sar** or **nsar** runs, it accumulates data on all system activities, sampling system data at regular intervals until it has produced a given number of samples. When you start data collection, you may specify the interval between samples and the number of samples. When you display the data, you may specify which types of data to review, or you may review all of it. For complete information on the system activity monitor, see **sar**(1) , **sar**(1M), **nsar**(1), and **nsar**(1M).

Starting System Activity Monitoring

Select the **sysadm** operation `System -> System Activity -> Start` to begin collecting data. The Start operation presents these prompts:

`Data Collection Name`
: Enter a file name to be used for the collection of data. If a data collection by the same name already exists, the monitor will append the new data to the existing file. If you do not specify a collection name, the operation will create a name of the form **spd.Day***nn*, where *nn* is the day of the month.

`Interval Between Samples (Seconds)`
: Enter the number of seconds that the system monitor should wait between taking samples. For best results, the interval should be not less than five seconds (the default) nor greater than a few minutes.

`Number of Samples`
: Enter the number of samples that you want the system monitor to take.

Sampling ends when the number of samples specified is reached. Keep in mind that each sample is over 1,000 bytes in size.

The system monitor stores the data in the file **/var/adm/sa/spd.***name*. The file remains until you delete it explicitly with the shell's **rm**(1) command or with the `System -> Activity -> Delete` operation.

You may run multiple system monitoring sessions at the same time.

Stopping System Activity Monitoring

Select the **sysadm** operation `System -> System Activity -> Stop` to halt a data collection session before the monitor has collected the specified number of samples. You may choose any existing data collections for the Stop operation.

Stopping a monitoring session does not remove the associated data file from **/var/adm/sa**. You may list the data from a prematurely—stopped monitoring session the same as for any monitoring session.

Deleting a System Activity Monitoring Data Set

Select the **sysadm** operation `System -> System Activity -> Delete` to remove one or more data collection files. Data collection files, located in **/var/adm/sa**, take up more than 1,000 bytes per sample, so you should delete them when no longer needed.

Displaying System Monitoring Data

Select the **sysadm** operation `System -> System Activity -> List` to review the data collected during a system monitoring session, select the List operation. You may list data from a monitoring session that is still in progress or from a session that has completed. If you list data from a monitoring session that is still in progress, the display includes all data collected up to that time.

The List operation lets you choose the kind of information you want to see in the report. Each category corresponds to a particular option of the **sar**(1) command, which is the program that collects the data. The data categories (and corresponding **sar** options, in parentheses) are:

`All data (A)`
: Data for all available categories.

`File access (a)`
: Use of system routines used for file access.

`I/O buffer activity (b)`
: Activity in system buffers, including hit ratios for system caches.

`System calls (c)`
: Number of system calls served for all system calls and for some specific system calls.

`Disk usage (d)`
: Physical disk I/O activity. Disk names displayed are long device specifications. To see how device specifications map to their entries in the **/dev** directory on your system, see the file **/etc/devlinktab**.

`Interprocess communications (m)`
: Interprocess message (**msgsnd**(2)) and semaphore (**semop**(2)) activity.

`Run queue and paging (q)`
: Number of processes running and waiting to run.

`CPU utilization (u)`
: CPU usage by user and system processes and idle time.

`Kernel tables (v)`
: Number of entries in the process table, inode table, file table, and shared memory record table.

`Paging I/O and process switches (w)`
: Process swapping and switching activity.

`Paging rates (p)`
: Paging activity such as virtual page faults and physical page faults.

`Free space (r)`
: Unused memory pages and swap area (disk) blocks.

`Terminal I/O activity (y)`
: Terminal (TTY) I/O activity.

For more information on **sar** output, see **sar**(1) and **sar**(1M).

Monitoring Process Activity

Select the **sysadm** `System -> Process` menu to monitor and control processes running on the system. A process is any program currently running on the system. The term *process* refers not only to the executing program code but also to the program's environment and state for that instance of execution.

The Process menu provides operations for deleting processes, changing process priority, listing processes, and sending signals to processes.

Deleting Processes

Select sysadm operation `System -> Process -> Delete` to terminate a process. Deleting a process removes it completely from memory and from the operating system tables. Deleting a process may also delete any of the process's child processes.

The Delete operation restricts the field of processes that you may delete according to criteria that you specify:

`Process ID`
: Enter the ID numbers of existing processes.

`Owner login name`
: Enter the login names of users whose processes you wish to delete.

`Terminal ports`
: Enter terminals (TTYs) whose associated processes you wish to delete; for example, **console**, **tty04**, **ttyp7**, and so on.

The operation then uses the **ps**(1) command to assemble a list of processes that satisfy the stated criteria. The operation prompts with `Process(es) to Delete`, letting you select processes from a list.

After the Delete operation has derived the desired process IDs based on your selections, it attempts to remove the processes with the equivalent of this command line:

```
# kill process_ID ↵
```

If this command does not succeed in killing them, it uses the equivalent of this command line:

```
# kill -9 process_ID ↵
```

Modifying Processes

Select the **sysadm** operation `System -> Process -> Modify` to change the priority of a running process. The priority of a process determines in part how much CPU time the process scheduler gives the process.

Every process has a priority value associated with it when it starts. In a possible range of 0 to 39, 0 is high priority and 39 is low priority. A process with a low priority number will tend to get more CPU time than a process of a higher priority number. The normal user process has a priority of 20.

You can alter processes' priority levels to reflect the importance of the jobs on your system. For example, if the daily backup process is competing with an urgent batch job, you can lower the priority (raise the priority number) of the backup process (the **dump2**(1M) command) to improve performance of the batch job.

The Modify operation restricts the field of processes that you may modify according to criteria that you specify:

`Process ID`
: Enter the ID numbers of existing processes.

`Owner login name`
: Enter the login names of users whose processes you wish to modify.

`Terminal ports`
: Enter terminals (TTYs) whose associated processes you wish to modify; for example, **console**, **tty04**, **ttyp7**, and so on.

The operation then uses the **ps**(1) command to assemble a list of processes that satisfy the stated criteria. The operation prompts with `Process(es) to Modify`, letting you select processes from a list.

You then select a new priority in the range 0 to 39, and the operation changes the process priorities with an equivalent of the **renice**(1M) command.

Displaying Processes

Select the **sysadm** operation `System -> Process -> List` to display a list of processes currently executing on the system. The List operation calls the **ps**(1) command to get the process status information. The operation lets you restrict the report to selected processes, accepting three kinds of selection criteria:

`Process ID`
: Enter the ID numbers of existing processes.

`Owner login name`
: Enter the login names of users whose processes you wish to list.

`Terminal ports`
: Enter the names of terminals with which processes are associated, for example, **console**, **tty04**, **ttyp7**, and so on.

You may select whether to display a long listing or the default listing. The default listing corresponds to the **ps** command's **-f** (full) listing. The long listing corresponds to the **ps** command's **-l** (long) listing. See the **ps**(1) manual page for more information.

Signaling Processes

Select the **sysadm** operation `System -> Process -> Signal` to send a signal to one or more processes. The effect of the signal depends on the receiving process. To kill a process, use `System -> Process -> Delete`, or use the Signal operation to send signal 15. If signal 15 does not kill the process, use signal 9. For more information on signals, see the manual page for the **kill**(2) system call.

The Signal operation restricts the field of processes for the operation based on criteria that you specify:

`Process ID`
: Enter the ID numbers of existing processes.

`Owner login name`
: Enter the login names of users whose processes you wish to signal.

`Terminal ports`
: Enter terminals (TTYs) whose associated processes you wish to signal; for example, **console**, **tty04**, **ttyp7**, and so on.

The operation then uses the **ps**(1) command to assemble a list of processes that satisfy the stated criteria. The operation prompts with `Process(es) to Signal`, letting you select processes from a list.

After you have selected or entered the desired signal, the Modify operation sends the signal to the processes.

Building and Booting a Kernel

The kernel is the executable program that provides operating system services to all other programs running on the system. The kernel runs directly on the hardware, managing access to peripherals such as tapes, terminals, and disks, as well as handling requests from users and application programs. By default, your system's current kernel is the file **/dgux**.

Although the DG/UX system ships with a starter kernel, you need to build your own custom kernel to serve the specific needs of your system and users. From time to time, you may also need to rebuild your kernel to tune performance or to accommodate changes in the hardware or software configuration.

To help you manage your system's kernel, the System menu includes the Kernel menu, which provides these operations:

`Auto Configure`
: Build a kernel that includes default parameter assignments and support for all hardware devices installed at standard locations.

`Build`
: Build a kernel that is the same as the one built with Auto Configure except that you can edit the system file and customize the tunable parameter assignments and add entries for nonstandard devices.

`Reboot`
: Shut down the system completely (except for the hardware itself) and restart the operating system.

The following sections elaborate on these operations.

Building a Kernel

You may select either of two sysadm operations to build a kernel: `System -> Kernel -> Auto Configure` **or** `System -> Kernel -> Build`. **These two operations are fundamentally the same except that the Build operation lets you edit the system file before continuing with the build process. There are also a few minor differences.**

In general, building a kernel involves creating a system file to reflect your hardware and software configuration. The Build or Auto Configure operation then uses the system file to determine what functionality to include in the kernel.

The system file can contain parameters that tune the operation and performance of your system. If you select the Build operation, you can add, remove, and change these parameters as you edit the system file. For more information about tunable parameters, see "Tuning System Parameters" at the end of this chapter.

Besides setting tunable parameters, the primary reason for editing the system file is to verify that the devices listed there reflect the devices and device drivers installed on your system. To make this task easier, the DG/UX system offers an autosizer, called **probedev**(1M), which looks for standard devices on your system and produces a list of any that it finds in standard locations. When the Auto Configure or Build operation creates a new system file, it uses **probedev** to generate the list of installed standard devices for inclusion at the beginning of the file.

If your system has nonstandard devices or device drivers, you need to add them to the list in the system file. For a complete list of devices that **probedev** recognizes, see **/usr/etc/probedevtab**.

If you do not wish to tune your kernel at this time, and if you do not have nonstandard devices installed on your system, you will find that the Auto Configure operation is the easiest way to build a kernel. If you have nonstandard devices on your system or if you want to change the defaults of any tunable parameters, you should use Build instead.

You should also use Build if one of the following conditions is true:

- Your system is an OS server, and you want to build a kernel for an OS client. Typically, OS servers and OS clients have different devices.
- Your system is an OS client that has local devices (such as disk or tape).
- Your system is one of the OS client hybrid configurations. An OS client hybrid configuration is one where **root** and/or **swap** reside on a local disk while the rest of the DG/UX file systems reside on another host on the network. Hybrid OS client configurations require some extra setting up, not covered in this manual. If your system is one of these hybrid configurations, see *Customizing the DG/UX™ System* for more information.

Whether you build a kernel now with Auto Configure or with Build, you may build a new kernel with either operation at any time.

Building a Kernel with Auto Configure

The Auto Configure operation starts by presenting you with these queries:

If your system has exclusive access to a disk–array storage subsystem (that is, your host is the only one using disks in such a subsystem), you may find the Auto Configure operation helpful for building your kernel. If two hosts use disks in the same disk–array subsystem, however, you should use Build to build your kernel.

`System configuration file name`

This name distinguishes the system file and the kernel from existing system files and kernels. The operation names the system file **system.***name* and places it in **/usr/src/uts/aviion/Build**, which is a symbolic link to **/var/Build**. The operation also uses this name when it creates the kernel, naming it **/dgux.***name*. You may enter a new name, or you may select the name of an existing system file.

If you select the name of an existing system file, the operation asks if you want to overwrite it by creating a new system file. If you elect to overwrite the system file, the operation also overwrites the associated kernel with the new kernel.

If you choose not to overwrite the system file, you return to the **System configuration file name** prompt.

`Operating system (OS) client`

Select this attribute if your system is a typical OS client of another host. A typical OS client has these characteristics:

- It boots the network device, **inen()**, instead of a disk.
- It mounts swap space from another host on the network.
- It mounts **root** and **usr** file systems from another host on the network.

Do not select this attribute if your system has the operating system installed on its own disk and does not depend on another host for the services described above.

After confirming the system file name that you have specified, the operation proceeds to assemble a new system file by concatenating a list of installed standard devices (generated by **probedev**(1M)) and the prototype system files supplied with the DG/UX system and other installed packages. The operation then builds the kernel. The following section, "Building a Kernel with Build," describes the system file and the build process in more detail.

When the build is complete, the new kernel is **/dgux.***name*, where *name* is the system file name you selected at the beginning of the operation. The kernel file **/dgux.***name* is linked to **/dgux**, which is the file that your system boots.

The new kernel will be in effect the next time you boot the system.

If you install new hardware or software on your system, you may have to build a new kernel. The kernel recognizes only those devices listed in the system file used to build the kernel.

Building a Kernel with Build

The Build operation is similar to the Auto Configure operation except in the following two ways:

- If the system file does not already exist, Build uses **probedev**(1M) to build a list of devices. The list is the same as the one that Auto Configure would compile if executed.
- Build lets you edit the system file before building the kernel.

The Build operation presents the following queries:

`System configuration file name`

This name distinguishes the system file and the kernel from existing system files and kernels. The system file is **system.***name*, located in **/usr/src/uts/aviion/Build**, which is a symbolic link to **/var/Build**. The operation also uses this name when it creates the kernel, naming it **/dgux.***name*. You may enter a new name, or you may select the name of an existing system file.

If you select the name of an existing system file, the operation lets you edit the file before performing the build. The operation does not generate a new

list of installed devices for the system file; therefore, if you have installed new devices, you have to add the entries to the file yourself.

If you enter a new name for a system file, the operation creates it by:

1. Invoking **probedev** to generate a list of installed standard devices. The list goes at the beginning of the file.
2. Concatenating the existing prototype system files, named **system.***package***.proto** in **/usr/src/uts/aviion/cf**, to form the body of the new system file.

`Build for this host or for OS client(s) of this host`
Select `this host` for the normal case, where you are building a kernel for this system.

Select `OS client(s)` if you are building a kernel for use by one or more OS clients. The remainder of this section describes the implications of these options in more detail.

If your system is a hybrid system, one that has a local disk but gets all or part of its operating system services from an OS server host, see *Customizing the DG/UX™ System* for more information.

`Boot Device: [inen(0)]`
Specify the network controller boot device that the OS client will use, or take the default if the boot device is **inen(0)**. Instead of an integrated network controller (**inen**), you can specify a Data General second generation network controller, name **dgen(***n***)**, where *n* is a number 0 through 5.

`Editor`
Enter the full pathname of the editor that you wish to use. The default is **vi**(1).

Next, the operation starts the selected editor and lets you edit the system file. The system file consists of prototype files from the various software packages installed on your system, concatenated to form one large file.

The system file contains entries for hardware devices, configuration variables, pseudo-devices, protocols, streams modules, tunable system parameters, and so on, as determined by the needs of the installed software packages. The file contains comments to help you understand the contents.

The system file entries of particular interest to you are:

Hardware devices
If you are building the kernel for this host, Build uses **probedev** to produce a list of currently-configured devices. This list appears at the top of the system file. If you are building the kernel for an OS client, Build instead inserts a standard list of devices typically found on OS clients. This list appears near the beginning of the file, after some parameter settings. You should review the list of devices to verify that it reflects the configuration for which you are building the kernel.

For reference, the system file already contains entries (preceded by comment characters) for some standard devices. These entries appear in the system file under this heading:

```
##### Typical AViiON OS client device configuration
```

The complete list of standard devices in standard locations is in **/usr/etc/probedevtab**.

If you intend to install additional SCSI disks or tapes on your system in the future, you can avoid having to rebuild the kernel by using abbreviated SCSI device specifications in the system file. Instead of specifying the SCSI ID in the device specification, use an asterisk (*). The asterisk represents all SCSI IDs on that controller. For example, instead of including these device specifications:

```
sd(insc(),0)
sd(insc(),1)
sd(insc(),2)
```

simply include this specification instead:

```
sd(insc(),*)
```

A kernel built with this specification will recognize a SCSI device at any SCSI ID on that specific controller. In general, you may include any of the following abbreviated SCSI device specifications (as appropriate for your system) in the system file:

```
sd(insc(),*)
st(insc(),*)
sd(cisc(),*)
st(cisc(),*)
sd(ncsc(),*)
st(ncsc(),*)
sd(dgsc(),*)
st(dgsc(),*)
da(hada(),*)
```

CAUTION: If building a kernel for a system in a dual-initiated configuration (one where two systems share a SCSI bus), do not use the asterisk notation described above. If two systems sharing a SCSI bus configure the same device at boot, a SCSI bus race condition will occur, and neither system will be able to access any devices on the bus.

If you are building the kernel for an OS client, the system file contains three tunable parameters at the beginning. These parameters are necessary for the system to function as an OS client. If you are building the kernel for this host, the system file does not contain these tunable parameters. If you want this host to function as an OS client, you must add these parameter declarations:

```
NETBOOTDEV          "inen()"
ROOTFSTYPE          NETWORK_ROOT
SWAPDEVTYPE         NETWORK_SWAP
```

Other than these three parameters, the tunable parameters are the same whether building for an OS client or for this host.

To review the entire set of available parameters, see the text files in **/usr/etc/master.d**. This information determines what devices the DG/UX system can recognize and what values a number of system parameters will have. The information for the DG/UX system itself is in the file **dgux**. If you have additional products, such as TCP/IP or ONC/NFS, their configuration files are in this directory also.

NOTE: You may view the master files in **/usr/etc/master.d**, but do not edit them. To override any settings in a master file, add the appropriate entries to your system file. Do not duplicate any of the master files in the master directory, or you will not be able to build your kernel.

For more information on tuning parameters, including a discussion of the parameters that you are most likely to want to change, see "Tuning System Parameters," later in the chapter.

After you edit the system file, the next step depends on whether you are building the kernel for an OS client of this host or for this host. If you are building the kernel for **OS client(s)**, the operation now tries to build the kernel. If you are building the kernel for **this host**, the operation presents the following prompt.

`Link the new kernel to /dgux`

After building the kernel executable, the operation moves it to the root directory as **/dgux.***name*. If you choose to link the new kernel to **/dgux**, the operation creates a link (using the **ln**(1) command) from **/dgux** to the new kernel so that the new kernel will be the one that boots when you next boot the system.

If you choose not to create the link, the existing kernel remains linked to **/dgux**.

To build the kernel, the operation first runs **config**(1M) on the system file and produces program code in a file named **conf.c**. If **config** fails, see "Configuration Error Messages." Correct the problem and invoke Build again. After **config** succeeds, the operation compiles **conf.c** and links the libraries in **/usr/src/uts/aviion/lb** to build the new kernel image. After successful completion, the bootable kernel file is in either of two places:

/ If you built the kernel for **this host**, the kernel is in the root directory, as **dgux.***name*. If so directed, the operation also linked the kernel to **/dgux**.

/usr/src/uts/aviion/Build

If you built the kernel for **OS client(s)**, the kernel is in this directory, which is a link to **/var/Build**. The kernel is named **dgux.***name*. You should move the kernel to some directory that is accessible to the OS clients.

Typically, OS client kernels reside in **/srv/release/PRIMARY/root/_Kernels**, or the equivalent directory for a

secondary release. Assuming that the OS client root directories are in the same file system as the **_Kernels** directory, you can now create links (using **ln**(1)) from **dgux** in the OS client root directories to the new kernel in **_Kernels**. For example, after moving OS client kernel **dgux.diskless** to the **_Kernels** directory, type the following to make it available to OS client **goober**:

```
# cd /srv/release/PRIMARY/root/goober ↵
# ln ../_Kernels/dgux.diskless dgux ↵
```

A new kernel takes effect the next time the system boots.

For more information on installing and setting up OS clients, see *Customizing the DG/UX™ System*.

If you install new hardware or software on your system, you may have to build a new kernel. The kernel recognizes only those devices listed in the system file used to build the kernel.

Configuration Error Messages

The following error messages are generated by the **config**(1M) program. Some errors originate in the master file, others in the system file. Errors in the system file are more common since you change it as a result of updating your configuration. Errors in the master file are less common; normally you do not alter the master file. You alter the master file only if you install a new device driver.

`A master file entry for` *`entry`* `already exists.`

Either you edited a master file and in doing so duplicated an entry in it, or you duplicated an entire master file. Make sure **/usr/etc/master.d** contains only the original, unchanged master files that you received with your software.

`Cannot open the master file [`*`master_file_name`*`].`
`Cannot open master file directory.`

Cannot open a file or directory. Make sure the master file is in the proper directory and that it is named correctly.

`No section definition found in master file [`*`master_file_name`*`].`
`This file will be ignored.`

The file in the master directory is not a legal master file.

`Unknown Keyword: [`*`keyword_name`*`]`
`Unknown Device Flag: [`*`device_flag`*`]`
`Unknown Flag: [`*`flag`*`]`

There may be incorrect information in your system file, such as a misspelled device name. Check that entries in your system file match those in your master file. Keyword errors pertain to the system file. Device flag and flag errors pertain to the master file.

```
Cannot Allocate Space.
Allocate device entry: Out of memory.
Allocate stream entry: Out of memory.
Allocate protocol entry: Out of memory.
Cannot allocate an alias structure.
Cannot allocate a keyword structure.
Error allocating Configured Device entry: Out of memory.
```

Cannot allocate space for internal structures. This is the result of an error returned from **malloc**(3C). This is related to user logical address space. Check the master file directory for duplicate files.

```
Illegal Master file line: [line]
Illegal protocol number: [protocol number]
Illegal Domain number: [domain number]
Illegal Socket number: [socket number]
Illegal Device Code: [device code]
No value associated with the keyword: [keyword_name].  Keyword
  will be ignored.
```

Illegal format for a master file or system file line. Device code errors and keyword errors are associated with the system file. The other errors in this category are associated with the master file.

Booting a Kernel

Select the **sysadm** operation `System -> Kernel -> Reboot` to boot your kernel. In the following example, the kernel **/dgux** on the **root** virtual disk located on **sd(insc(0),0,0)** is booted automatically to run level 3. You could specify a different run level.

```
Boot path: [sd(insc(0),0,0)root:/dgux -3] ↵
All currently running processes will be killed. ↵
Are you sure you want to reboot the system? [yes] ↵
```

The system is shut down and rebooted to run level 3. See Chapter 3 for more information about run levels and booting your system.

Setting and Displaying DG/UX Parameters

The DG/UX parameters control what actions occur on your system when it boots. These actions include checking file integrity, cleaning up after unfinished jobs of various kinds, and starting programs that provide various system services. The parameters also control a number of other things.

To display these parameters, which appear in **/etc/dgux.params**, use the **sysadm** operation `System -> Parameters -> Get`. To turn boot actions on or off, use the operation `System -> Parameters -> Set`.

A list of the actions available through the Set operation appears below, followed by a list of DG/UX parameters not accessible through the Set operation. There is a help message for each parameter in the Set operation.

`Print Verbose Messages at Boot`
Select this feature if you want verbose messages during each run level change. The verbose messages are always written to the **/etc/log/init.log** file. If you do not want to see the verbose messages, leave this feature set to its default value, off.

`Check Password File at Boot`
Select this feature if you want the system to check the password file, **/etc/passwd**, for entries that lack passwords each time the system boots. Leave this feature unselected if you do not want the system to look for profiles without passwords.

`Check UUCP Files at Boot`
Select this feature if you want the system to verify that UUCP file permissions are correct each time the system boots. The system corrects any inappropriate permissions. If you do not use UUCP or if you are confident that the file permissions are correct, do not select this feature.

`Check for Packages Needing Setup at Boot`
Select this feature if you want the system to check at boot time for software packages that you need to set up. You may want to use this feature if you frequently add new software packages. If you seldom add new software packages, you may not want to select this feature. If you do not select this feature, the system will not alert you that you have to set up new packages.

`Start File System Checker Without Verification at Boot`
Select this feature if you want the system to start checking the integrity of file systems without querying you each time the system boots. If you do not select this feature, the system will ask you every time it boots if you want to check file systems. If you do not select this feature, realize that the boot process will not continue until you have entered a response at the system console. The system uses **fsck**(1M) to check the integrity of file systems.

`Boot without Verifying Date and Time`
Select this feature if you want the system to boot without asking you to verify the date and time. Leave this feature unselected if you want the system to pause during boot and ask you to verify the system date and time. If you do not select this feature, realize that the boot process will not continue until you have typed some response at the system console. You may use the **date**(1) command to set the date and time while the system is running.

NOTE: Setting the date and time while the system is at run level 1 or higher may cause some system services (or daemons) to behave improperly. For best results, take your system to single-user mode before setting date and time.

`Download Async Ports at Boot`
Select this feature to load and start asynchronous controllers at boot. If you do not select this feature, your system terminals (and any other asynchronous devices, for example, some printers) will not be available when the system comes up. The system uses the **tcload**(1M) command to load asynchronous controllers.

`Mount Local File Systems at Boot`
Select this feature to mount all local file systems at boot. Leave this feature unselected if you have no local file systems or if you do not want them mounted at boot. Local file systems are the ones listed in **/etc/fstab** that are not of type **nfs**. Mounting a file system makes it available for use on the system.

`Start wmtd Daemon at Boot`
Select this feature to start the **wmtd**(1M) daemon at boot. The **wmtd** daemon allows you to use a WORM (write once, read many) device as a tape device. If you have a WORM device, you may want to select this feature. If you want to specify any arguments for the **wmtd** command line, specify them at the `Arguments to wmtd` prompt in this operation. Leave this feature unselected if you do not have a WORM device.

`Start System Error Log Daemon at Boot`
Select this feature to start the system error log daemon, **syslogd**(1M), at boot. We recommend that you start **syslogd** at boot because some system facilities (such as the disk mirroring portion of the kernel) use it. For more information on the system error logger, see Chapter 16.

`Start System Accounting at Boot`
Select this feature to start system accounting each time the system boots. Make sure the appropriate **cron** job entries appear in the superuser's **crontab** file. See Chapter 14 for more information on the accounting system. See "Automating Job Execution" in Chapter 2 for more information on **cron**.

`Start cron Daemon at Boot`
Select this feature to start the **cron** daemon each time the system is booted. You should run this daemon unless you are sure that none of your system's users or packages are using this service. This service is useful for running commands submitted either via a **crontab** file or with the **at** or **batch** commands. If you use the accounting system, the LP (printer) subsystem, or the UUCP facility, you must also run the **cron** daemon. See "Automating Job Execution" in Chapter 2 and the **cron**(1M) manual page for more information on **cron**.

`Start lpsched Printer Scheduler at Boot`
Select this feature to start the **lpsched**(1M) printer scheduler daemon each time the system is booted. This daemon must be running in order for users to print jobs with the **lp**(1) command. This daemon also supports print requests submitted with the **lpr**(1) command. If not started at boot time, you can start this daemon later with **sysadm**(1M).

`Start lpd Printer Scheduler at Boot`
Select this feature to start the **lpd**(1M) printer scheduler daemon each time the system is booted. This daemon supports print requests submitted with the **lpr**(1) command.

`Preserve Editor Temporary Files at Boot`
Select this feature to run the **expreserve** daemon each time the system is booted. **expreserve** saves the temporary files that are left behind when a system crash or other interruption causes **ex**(1) or **vi**(1) editing sessions to

quit prematurely. If necessary, the **expreserve** daemon sends mail to users telling them how to restore interrupted editing sessions. If any users on your system use **ex** or **vi**, you should run this daemon.

`Number of biod daemons`

Enter the number of **biod** daemons to be started when the system boots. **biod** daemons are used in performing asynchronous I/O between secondary storage and main memory except for paging to local swap areas. For example, file readahead, most file buffer writebacks, and paging to a remote disk all use **biod** daemons. The more I/O expected on the system, particularly ONC/NFS I/O, the more daemons are needed to service it. A good value for a typical system using ONC/NFS is 8; fewer **biod** daemons may yield equally good performance if the system is not used as an ONC/NFS client (that is, if it does not access remote file systems much).

`Arguments to swapon`

Enter arguments to the **swapon**(1M) command. The **swapon** command runs each time you boot the system, checking the **/etc/fstab** file for entries of type **swap** and making those virtual disks available to the system as additional paging area. You may specify the **–a** option to specify that the system should use all swap areas appearing in **/etc/fstab**, or you may specify particular virtual disks (for example, **/dev/dsk/swap_alt**). Normally, the argument is a null string, which specifies the one default swap device.

`Arguments to wmtd`

Enter the virtual-to-physical device mapping used by the WORM-as-magnetic-tape server daemon. For example, if you want **/dev/wmt/0** and **/dev/wmt/0n** to be associated with the device **/dev/rpdsk/2**, then enter **0=/dev/rpdsk/2**. See **wmtd**(1M) for more information. Refer to the file **/etc/devlinktab** to see how virtual devices map to physical devices on your system. If you do not have a WORM device on your system, you may ignore this query.

`Arguments to expreserve`

Enter the names of the directories where the **vi**(1) and **ex**(1) editors create temporary files during editing sessions. The **expreserve** daemon checks these directories for files left by prematurely terminated **ex** and **vi** sessions. **expreserve** moves these abandoned temporary files to **/var/preserve**.

Some DG/UX parameters are not accessible through the Set operation. To change these, edit the **/etc/dgux.params** file. The parameters are:

dkctl_START

This parameter determines whether or not the system will enable the write-verify mode of operation for selected disks. Select disks for write-verify operation with the **dkctl**(1M) command. Once set, write-verify defaults appear in the **/etc/default/dkctl** file. The default is **true**. For more information on the write-verify feature, see Chapter 7.

fsck_ARG

This parameter provides the arguments to be supplied to the file system checker, **fsck**(1M), when it runs at boot time. The default value is **–xlp**.

reboot_notify_START
: This parameter determines whether or not the system sends notification mail when it reboots. Upon rebooting, the system sends mail to any local users listed in **reboot_notify_ARG**, below. The default is **false**.

reboot_notify_ARG
: This parameter determines which local logins, if any, receive mail indicating that a system reboot has occurred. To receive notification, you must also set **reboot_notify_START**, above, to **true**. By default, the parameter is set to an empty value.

strtty_START
: This parameter determines whether or not the system pushes STREAMS modules for terminal devices at boot. The default is **true**.

Other parameters also appear in **/etc/dgux.params**. To change them, edit the file. The file contains comments to help you understand the parameters and the accepted values.

Setting Global Locale Variables

The **sysadm** `System -> Language` menu provides operations for setting and listing the native language support parameters for your system.

The native language parameter, **LANG**, tells the system which locale database to use for determining functionality that varies from region to region throughout the world. For example, the locale database determines currency symbols, collating sequences for sorting operations, comma and decimal point usage and placement for numerals, and so on. Certain DG/UX commands, including **cp**(1), **find**(1), **ln**(1), **mv**(1), **rm**(1), and **tar**(1), accept native language responses to yes/no questions. For example, if you set the **LANG** variable to a French language locale, these commands accept **oui** and **o** in addition to **yes** and **y**.

The native language support facility also provides support for applications that use the X/Open message facility message catalogs. A message catalog is a list of strings in a given language intended for use in a specific program or application. Some applications ship with multiple message catalogs, each one supporting the application in a different language. The **NLSPATH** environment variable tells the system what directories to search for message catalogs.

By default, the **NLSPATH** variable includes directories based on the value of the **LANG** variable. If you install software that loads message catalogs into other directories, you need to add these directories to **NLSPATH**.

Use the Set operation to set your system's global **LANG** and **NLSPATH** variables. Use the Get operation to display the variables' current values.

DG/UX 5.4R3.00 includes locale databases for 70 different locales. A locale does not necessarily correspond to a single language, country, or geographic region. Locale names, which are the values you may assign to the **LANG** variable, have the form

language [*_territory* [*.codeset*]]

Some examples are **fr** (French), **fr_CA** (Canadian French), and **fr_CA.850** (Canadian French, using the IBM 850 code set). The code set normally available on DG/UX system consoles and X Window System windows is ISO 8859-1, the Western European code set. Data General terminals and printers support either ISO 8859-1 (or a very close approximation of it) or ASCII. The ISO 8859-1 codeset is a superset of ASCII. The locales that specify other code sets are useful only with I/O devices (terminals and printers) that use those code sets.

The locale definitions supplied as part of the DG/UX system appear in the following table. These are the values that you may assign to the **LANG** variable. These all appear as subdirectories of **/usr/lib/locale**. You are free to delete any that you do not need; however, you must leave the **C** locale definitions intact. The contents of the locale definitions are described in the **setlocale**(3C), **colltbl**(1M), **chrtbl**(1M), **montbl**(1M), and **mkmsgs**(1) manual pages.

Table 4–2 shows the locale names and related information for locales supported on the DG/UX system.

Table 4–2 Supported Locales

Locale Name	Language	Country	Code Set	Equivalent Name
C	English	US	ASCII	C–locale
sk	Czech	Czechoslovakia	ISO 8859–2	
da	Danish	Denmark	ISO 8859–1	da_DK
da_DK.850	Danish	Denmark	PC 850	
da_DK.865	Danish	Denmark	PC 865	
nl	Dutch	Netherlands	ISO 8859–1	nl_NL
nl_NL.437	Dutch	Netherlands	PC 437	
nl_NL.850	Dutch	Netherlands	PC 850	
nl_BE	Dutch	Belgium	ISO 8859–1	
nl_BE.437	Dutch	Belgium	PC 437	
nl_BE.850	Dutch	Belgium	PC 850	
en	English	United Kingdom	ISO 8859–1	en_GB
en_GB.437	English	United Kingdom	PC 437	
en_GB.646	English	United Kingdom	ISO 646	
en_GB.850	English	United Kingdom	PC 850	
en_AU	English	Australia	ISO 8859–1	
en_AU.437	English	Australia	PC 437	
en_AU.646	English	Australia	ISO 646	
en_AU.850	English	Australia	PC 850	
en_CA	English	Canada	ISO 8859–1	
en_CA.437	English	Canada	PC 437	
en_CA.646	English	Canada	ISO 646	
en_CA.850	English	Canada	PC 850	

Continued

Table 4–2 Supported Locales

Locale Name	Language	Country	Code Set	Equivalent Name
en_US	English	United States	ISO 8859–1	
en_US.437	English	United States	PC 437	
en_US.646	English	United States	ISO 646	
en_US.850	English	United States	PC 850	
fi	Finnish	Finland	ISO 8859–1	fi_FI
fi_FI.437	Finnish	Finland	PC 437	
fi_FI.850	Finnish	Finland	PC 850	
fr	French	France	ISO 8859–1	fr_FR
fr_FR.437	French	France	PC 437	
fr_FR.850	French	France	PC 850	
fr_BE	French	Belgium	ISO 8859–1	
fr_BE.437	French	Belgium	PC 437	
fr_BE.850	French	Belgium	PC 850	
fr_CA	French	Canada	ISO 8859–1	
fr_CA.850	French	Canada	PC 850	
fr_CA.863	French	Canada	PC 863	
fr_CH	French	Switzerland	ISO 8859–1	
fr_CH.437	French	Switzerland	PC 437	
fr_CH.850	French	Switzerland	PC 850	
de	German	Germany	ISO 8859–1	de_DE
de_DE.437	German	Germany	PC 437	
de_DE.850	German	Germany	PC 850	
de_AT	German	Austria	ISO 8859–1	
de_AT.437	German	Austria	PC 437	
de_AT.850	German	Austria	PC 850	
de_CH	German	Switzerland	ISO 8859–1	
de_CH.437	German	Switzerland	PC 437	
de_CH.850	German	Switzerland	PC 850	
el	Greek	Greece	ISO 8859–7	el_GR
is	Icelandic	Iceland	ISO 8859–1	is_IS
is_IS.850	Icelandic	Iceland	PC 850	
it	Italian	Italy	ISO 8859–1	it_IT
it_IT.437	Italian	Italy	PC 437	
it_IT.850	Italian	Italy	PC 850	
it_CH	Italian	Switzerland	ISO 8859–1	
it_CH.437	Italian	Switzerland	PC 437	

Continued

Table 4–2 Supported Locales

Locale Name	Language	Country	Code Set	Equivalent Name
it_CH.850	Italian	Switzerland	PC 850	
no	Norwegian	Norway	ISO 8859–1	no_NO
no_NO.850	Norwegian	Norway	PC 850	
no_NO.865	Norwegian	Norway	PC 865	
pl	Polish	Poland	ISO 8859–2	pl_PL
pt	Portuguese	Portugal	ISO 8859–1	pt_PT
pt_PT.850	Portuguese	Portugal	PC 850	
pt_PT.860	Portuguese	Portugal	PC 860	
ru	Russian	Russia	ISO 8859–5	ru_SU
sh	Serbo Croatian	Yugoslavia	ISO 8859–2	sh_YU
es	Spanish	Spain	ISO 8859–1	es_ES
es_ES.437	Spanish	Spain	PC 437	
es_ES.850	Spanish	Spain	PC 850	
sv	Swedish	Sweden	ISO 8859–1	sv_SE
sv_SE.437	Swedish	Sweden	PC 437	
sv_SE.850	Swedish	Sweden	PC 850	
tr	Turkish	Turkey	ISO 8859–3	

Setting Time and Date

To set the time and date, use the **sysadm** operation `System -> Date -> Set`. When you set the time, you may specify:

- Month
- Day of the Month
- Hour
- Minute
- Year
- Time Zone

For a time zone, select one of those listed on the screen or in the help message. See the **timezone**(4) manual page for more information on the time zone format.

We recommend that you do not set the time back while the system is at run level 1 or higher. Setting the time back while at any level above single-user mode may cause system daemons and X Window System clock clients to behave erratically. Take the system down to single-user mode before setting the time back.

You may set the time forward at any run level without disrupting the function of system daemons.

To display the time and date, use the **sysadm** operation `System -> Date -> Get`. In the shell, you set and display the date and time with the **date**(1) command.

Alternatively, you may prefer to get time from a LAN–based designated clock server, providing the precise time to all hosts on the LAN. Refer to *Managing TCP/IP on the DG/UX™ System* for information to set up NTP (Network Time Protocol) in your network.

Improving Performance

This section contains suggestions for improving the performance of your system. See *Analyzing DG/UX™ System Performance* for more information.

Also, see "Tuning System Parameters" later in this chapter for tunable parameter charts, definitions, and recommendations.

High Availability

High availability is the integrated set of system resources and services required to maximize system availability and minimize downtime. These resources and services combine to provide near–continuous applications availability at a reasonable cost. No single system component can provide high availability; however, the selection of hardware and software high availability elements can meet your particular productivity requirements.

Chapter 7 in this manual discusses the attributes of Online Storage Management (OSM), which enables the manipulation of data storage resources while the data is online and in use. OSM offers unprecedented performance and availability gains.

Refer to *Achieving High Availability on AViiON® System* for an introduction to the concept of high availability and Data General's solutions.

Maximizing System Usage

To ensure maximum system performance, you should check for:

- Less important jobs interfering with more important jobs
- Unnecessary jobs
- Jobs running during peak hours that could just as well run during off-peak hours
- The efficiency of user-defined features, such as personal configuration files (**.profile** and **.login**) and the PATH environment variable

Getting Process Information

To obtain information about active processes, use the **sysadm** operation `System -> Process -> List`. This operation calls the **ps**(1) command to produce a listing.

The listing constitutes a "snapshot" of what is going on, which is useful when you are trying to identify what processes are loading the system. Things will probably change by the time the output appears; however, the entries that you should be interested in are TIME (minutes and seconds of CPU time used by processes) and STIME (time when process first started).

If you spot a "runaway" process, one that uses progressively more system resources over a period of time while you are monitoring it, you should probably stop the process with the operation `System -> Process -> Delete`.

If you regularly run processes that take a very long time to execute, you should consider using **cron**(1M) or **at**(1) to execute the job during off-hours, or use **batch**(1) to execute the job when system load level permits. See "Automating Job Execution" in Chapter 2 for more information on these commands.

Checking User Search Path Variables

Every shell process has a **path** or **PATH** environment variable that lists the directories that the system should search when looking for a command invoked by the user. Every time the user issues a command, the system scans the directories in the path to see where the command resides. If the command invokes other commands, the system has to scan the search path for them too. These searches require both processor and disk time, thus changes here can help performance.

Some things that you should check for in user search path variables are:

Path Efficiency
: The system searches the path directories in the order listed, so your most commonly used directories should appear first in the path. Make sure that a directory is not searched more than once for a command.

Path Length
: In general, the search path should have the least number of required entries.

Large Directory Searches
: Avoid searching large directories if possible. Put any large directories at the end of the search path.

Convenience and Human Factors
: Users may prefer to have the current directory listed first in the path (**.:/usr/bin**), but note that putting the current directory first can lead to breaches in security. If a program having the same name as a system command, such as **ls** or **pwd**, for example, is in the user's current directory, the user will inadvertently execute it while intending to execute the "real" system command. Depending on the nature of the "fake" system command that exists in the local directory, the results could be undesirable, or at best, unpredictable.

  It is particularly critical that the superuser profiles (**root** and **sysadm**) not have the current directory first on their path. In fact, it is safer to leave the current directory out of a superuser path altogether.

Shift Workload to Off-Peak Hours

Use the **crontab** command to examine users' **crontab** files to see if there are jobs scheduled for peak hours that could just as well run during off hours. You may find the accounting system (Chapter 14) and the Process menu's List operation (which invokes the **ps**(1) command) helpful in determining what processes have the greatest impact on system performance.

Encourage users to run large, noninteractive jobs (such as program compilations) at off-peak hours. You may also want to run such jobs with a low priority by using the **nice**(1) or **batch**(1) commands. As superuser, you can always change a job's priority with the **renice**(1M) command.

Tuning System Parameters

Tunable system parameters set various table sizes and system thresholds to handle the expected load on your system. You'll find the default tunable parameter values are adequate for most configurations and applications. If your application has special performance needs, you may have to experiment with different combinations to find an optimal set. To set tunable parameters, edit the values in your system file when building a new kernel with the Build operation.

Uname Configuration Variables

The **uname**(1) command and the UUCP system use the Uname configuration variables. Each of these variables must contain a character string no longer than 256 characters, including the trailing null character. These parameter variables are listed in **/usr/etc/master.d/dgux**.

NODE The UUCP nodename of the system (**sales**, **sys31**, and so on). The default is **no_node**.

MACH The name of the system's underlying hardware. The default is **AViiON**.

SYS The name of the operating system. The default is **dgux**.

VER The version number of the operating system. The default is **generic**.

REL The number of the system's release. The default for this release is **5.4R3.00**.

Sysinfo Configuration Variable

The **sysinfo**(2) system call uses this configuration variable. The variable must be contain a character string no longer than 256 characters, including the trailing null character. This parameter variable is also listed in **/usr/etc/master.d/dgux**.

HW_PROVIDER
Name of the company that built the computer hardware. The default is **Data General**.

Setup and Initialization Configuration Variables

The setup and initialization configuration variables are also listed in **/usr/etc/master.d/dgux**. These variables set the system initialization parameters shown in the following list.

DST The type of Daylight Savings Time being used. The different types are defined in **/usr/include/sys/time.h**. The default is **1**.

TZ The time zone of your system in minutes west of Greenwich Mean Time (GMT). Set this according to your time zone. The default is USA Eastern time: **300**.

DUMP The name of the default system dump device in DG/UX common device specification format. This device is the default device used to do a system memory dump after a panic or a halt. The default is **st(insc(),4)**.

DEBUGGER
The kernel debugger to be used by your system. The default is the null debugger stub. The DG/UX Kernel Debugger can be specified simply by listing the keyword **DEBUGGER** in your system configuration file, which causes the implied value (**&deb_debugger_request**) to be used. See *Using the DG/UX™ Kernel Debugger* for details. The default is **&sc_null_debugger**.

DEBINITCMDS
A list of zero or more commands to be executed by the DG/UX Kernel Debugger (if present) before displaying the first debugger prompt. If the null debugger is used, this variable has no effect. Note that debugger commands must be separated and terminated by the C language New Line character, **\n**. The default is **mode er on\n**.

INIT The internal function that is executed upon system booting. The default function calls the program specified in the INITPATH parameter. Do not change this parameter. The default is **&init_run_sbin_init**.

INITPATH
The system initialization program that the kernel should execute as the very first user process on the system. The default is **/sbin/init**.

NETBOOTDEV
The network device to be used in the booting of an operating system client receiving some or all of its operating system resources over the network. The default is a null string. Only set this variable if network-assisted booting should take place.

NETSTART
The internal function that will start the network internal to the kernel. The default is **STUB_NET**, meaning the network is not started by the kernel at all. Use **REAL_NET** to indicate that the network will be started by the kernel.

ROOTFSTYPE
The internal function that will mount the root file system. The default is **LOCAL_ROOT**, meaning the root file system will be found on a local disk device. Use **NETWORK_SWAP** to indicate that the swap file will be NFS-mounted over the network, or **MEMORY_SWAP** if swapping will be done entirely in memory (that is, no swapping at all).

SWAPDEVTYPE

The internal function that will initialize the swap file. The default is **LOCAL_SWAP**, meaning the swap device will be mounted from a local disk device. Use **NETWORK_ROOT** to indicate that the root file system will be NFS-mounted over the network, or **MEMORY_ROOT** if the contents of the root will be bound into the kernel image itself.

STARTER

A boolean variable indicating whether or not the system will ask to configure additional devices upon booting. The default is **0** (FALSE). Use **1** for TRUE. We recommend that you use the default.

SHOWGOODCONFIGS

A boolean variable indicating whether or not the system will print a message for every successful attempt to configure a device during system initialization. Use **1** for true and **0** for false. The default is **0**.

SHOWBADCONFIGS

A boolean variable indicating whether or not the system will print a message for every unsuccessful attempt to configure a device during system initialization. Use **1** for true and **0** for false. The default is **1**.

RUNFSCK

A boolean variable indicating whether or not the kernel should run **fsck** to check the root file system before trying to mount it. Use **1** for true and **0** for false. The default is **1**.

FSCKFLAGS

The options to be used when invoking **fsck** to check the root file system before attempting to mount it. This parameter is only relevant when RUNFSCK is true. The default is **-xlq**.

ROOTLOGSIZE

An integer giving the number of blocks to be used for **fsck**'s fast recovery log on the root file system. The default is **0**, meaning the root will not be mounted in fast recovery mode. Specify a number greater than **0** for the root to be mounted in fast recovery mode. For more information on fast recovery **fsck**, see Chapter 9.

ROOTREADONLY

A boolean variable indicating whether the root file system should be mounted read-only, instead of the default read-write. Use **1** for true and **0** for false. The default is **0**.

NSTRDEMONS

The number of streams demons to start at initialization. The default is **0**, meaning one demon per every 2 JPs is run (rounded up so that there is 1 demon if 1 JP, 2 if 3 JPs, etc.) If **NSTRDEMONS** is greater than the number of active JPs, one demon per active JP is run.

CPU, Process, and Memory Configuration Variables

The CPU, process, and memory configuration variables are also listed in **/usr/etc/master.d/dgux**.

NPROC

The maximum number of processes the system can have at one time. For various sized systems use the following values: small, **96**; medium (default), **256**; and large, **512**. The overall number of processes needed depends on the number of terminal lines available, the number of processes spawned by each user, and the number of system processes and network daemons. If the maximum number of processes is used up, the **fork**(2) or **vfork**(2) system call will return the error EAGAIN.

NCPUS

The actual number of processors to run. If set to **0** (the default), all available CPUs will be used. Any other value specifies that number of CPUs to run. If the value specified is more or less than the number of CPUs present, a message to that effect is printed when the kernel is booted. Note that on a uniprocessor system, this parameter has no real effect since the one processor will always be run.

NLWP

The maximum number of user LWPs the system can have at any one time. If set to **0** (the default), this value will be dynamically calculated based on the amount of available memory in the system.

NLWPGROUPS

The maximum number of user LWP groups that the system can have at any one time. If set to **0** (the default), this value will be dynamically calculated based on the amount of available memory in the system. An LWP group is a set of locally scheduled LWPs from the same process that share the same accounting and global priority. Note that the DG/UX LWP group is different from a USL ES/MP LWP set, which is always composed of globally scheduled LWPs from potentially different processes. The DG/UX LWP group may be thought of as an ES/MP subset or a general ES/MP LWP.

MAXBOUND

The maximum number of user LWPs (light–weight processes) that can be ready to run in the kernel, which is the same as the maximum number of bound transient data sections (i.e., wired kernel stacks). If set to **0** (the default), then the number of bound transients will be determined dynamically based on system load. This configuration parameter corresponds most closely to NVPS in previous releases. The maximum bound can never be less than the minimum bound. If MINBOUND is not **0** and MAXBOUND is not **0**, the the maximum bound will be the larger of MINBOUND and MAXBOUND.

MAXSLICE

The default dispatcher round–robin timeslice used for the SCHED_OTHER (timesharing) and SCHED_RR scheduling policies. If a user LWP/LWP group runs for this amount of time, it will yield the CPU to other LWPs/LWP groups with the same priority. Note that the round–robin timeslice is essentially infinite for SCHED_FIFO and SCHED_DG_LIFO scheduling policies. The default is **100** milliseconds.

MAXULWP

The maximum number of LWPs that a non–superuser can have in existence at any time. By default, there is no per–user limit.

MAXULWPGROUPS

The maximum number of LWP groups that a non–superuser can have in existence at any one time. By default, there is no per–user limit.

MAXLATENCY

The maximum time in milliseconds a user process can run before being interrupted to check whether it should be preempted by a process added to the eligible list by another JP. The default is **50**.

MAXUP

The maximum number of processes that a nonsuperuser can have in existence at one time. The default is **50**. This value should not exceed the value of NPROC (NPROC should be at least 10% more than MAXUP). This value is per user identification number, not per terminal. For example, if ten people logged in with the same user ID, the default limit would be reached very quickly.

SDESLIM

The default number of descriptors (soft limit) a process is allowed to have at one time. The default is **64**.

HDESLIM

The maximum number of descriptors (hard limit) a non-superuser process is ever allowed to have at one time. A non-superuser process may change its soft limit up to the value of the hard limit. The default is **1024**.

PERCENTLOCKABLE

The percentage of system memory available for locking by user processes, provided through the **memcntl**(2) system call. The default is **10**. Note that reserving and then locking large amounts of memory may deadlock your system.

MAXPAGEOUTS

The maximum number of pageout I/O operations that the system can have at one time. The default is **0**, indicating that the system should pick a reasonable value based on the amount of physical memory present on the system. See the entry in **/usr/etc/master.d/dgux** for a discussion of the implications of this for your system.

MAXDRIVERS

Specifies the number of slots that will be present in the device driver table. The default is **64**.

Pseudo-Device Unit Count Variables

The pseudo-device unit count variables set the number of units a specified pseudo-device will have. (No count variables are needed for real devices; any units present are usable.)

PTYCOUNT

The number of pseudo-terminal device pairs (**/dev/pts/***) that will be created when the system is booted. The default value is **64**. This parameter is used for **telnet**(1C), **rlogin**(1C), and **shl**(1).

PMTCOUNT

The number of pseudo-magnetic tape devices used for access to remote magnetic tapes (**/dev/pmt/***). The default is **20**.

WMTCOUNT

The number of pseudo-magnetic tape devices used for access to WORM (write-once/read-many) devices (**/dev/wmt/***). The default is **8**.

File System Configuration Variables

The file system configuration variables are also listed in **/usr/etc/master.d/dgux**. These variables set the file system parameters shown in the list below.

ACCTON, ACCTOFF

If the free space in the file system in which the accounting file resides becomes less than the percentage specified by ACCTOFF, no further accounting records will be written. When the free space reaches ACCTON percent, the writing of accounting records will resume. ACCTOFF should always be smaller than ACCTON. The default for ACCTON is **5**. The default for ACCTOFF is **2**.

MAXBUFAGE

The maximum amount of time that modified system data will remain in a buffer cache before being written to storage. The default is **60** seconds. A relatively high value keeps the data in the cache for a longer time period; a low value keeps the data in the cache for a shorter time period. Maximizing the amount of time that the data remains in the cache without flushing may produce improved performance. The tradeoff, however, is that less frequent flushes may contribute to unreliable data in the event of a system crash. If the system crashes, data remaining in the cache does not get written to storage.

MAXSYSBUFAGE

The maximum age in seconds that a modified buffer in main memory containing system data can reach before it is written to disk. If this parameter is **0**, the default, the system uses the same maximum age value for system buffers that it uses for user buffers (as set in MAXBUFAGE). If you set MAXSYSBUFAGE to a value greater than MAXBUFAGE, the system ignores it and uses MAXBUFAGE for all buffers.

PERCENTBUF

The maximum percentage of system memory that can be occupied by data files. The default is **100**, meaning data files can use as much of system memory as needed. There should be no need to configure this variable unless heavy data file usage is hampering system performance by provoking an undesirable amount of other paging.

PERCENTSYSBUF

The percentage of system memory (after initialization) that is reserved for system buffers. System buffers hold directory, inode, file index, and bitmap data. The default, **0**, causes the system to select a reasonable value for the system.

HOGFILESIZE

The maximum number of bytes of physical memory which can be used by a given data file before that file will be treated unfavorably for physical memory resource allocation. The default is **262144**. When the system is forced to page data out to meet requests for memory, data files having more than HOGFILESIZE bytes buffered are more likely to be paged out, and will be less able to consume free memory.

CDLIMIT

The maximum size in bytes that a nonsuperuser file may attain. The default is the constant **INT32_MAX**, which is equal to 2,147,483,648.

NFSLOCKUSERLIMIT

The maximum number of remote processes that can hold record locks concurrently. The default is **512**.

USERLOCKLIMIT

The maximum number of user locks that a process can hold. The default is **2048**.

FREEINODE

The maximum ratio of in-use inodes to free inodes in the system. To improve performance on systems where you open a large number of files repeatedly, set this parameter to a higher value. The default is **4**.

FREERNODE

The maximum ratio of in-use rnodes to free rnodes in the system. To improve performance on systems where you open a large number of files repeatedly, set this parameter to a higher value. The default is **4**.

NCLIENTOPS

The maximum number of handles of a given RPC type available on the system. The default is **6**.

NFS

The name of a table of file system functions which determines if NFS is present. The default is **&nfs_pkg_function_table_stub**, meaning that NFS is not present. The implied value is **&nfs_pkg_function_table_meat**, meaning that NFS is present.

CHOWN_REST

A boolean variable indicating whether or not the POSIX feature **_POSIX_CHOWN_RESTRICTED** is present for the system. The default is **0** (FALSE). Use **1** for TRUE. If the feature is present, then additional POSIX-style restrictions are placed on the **chown**(2) system call.

SRVNOTNEEDED

A boolean variable indicating whether or not the DG/UX file system should retain buffers read from NFS clients. The default is **0** (FALSE). If this value is **1** (TRUE), then the server will mark buffers read by NFS clients as unnecessary, which makes them more likely to be reused than other pages in the cache.

FULL_ISO9660

A boolean variable indicating whether or not Full ISO 9660 filenames will be used for files on High Sierra compact discs. Unless this behavior is turned on, the High Sierra file manager will map uppercase characters in filenames to lowercase, and map semicolon (;) characters in filenames to hyphen (-) characters. The default is **0** (FALSE). Use **1** for TRUE. We suggest you use the default.

STREAMS Configuration Variables

The STREAMS configuration variables are associated with STREAMS processing. We recommend that you use the default values supplied. These variables are also listed in **/usr/etc/master.d/dgux**. They set the STREAMS parameters shown in the following list.

PERCENTSTR

The percentage of system memory (after initialization) that is reserved for STREAMS buffers. The default is **20**.

NQUEUE

The maximum number of STREAMS queues (Streams plus instances of STREAMS modules) that may exist at any one time on the system. A minimal stream contains two queue pairs: one for the Stream head and one for the driver. Each instance of a module on a Stream requires an additional queue pair. The default is **2048**.

NSTRPUSH

The maximum number of STREAMS modules that may be pushed on any one Stream. This is used to prevent an errant user process from consuming all the available queue pairs on a single STREAMS module. The default is **9**.

STRMSGSZ

The maximum number of bytes allowed in the data portion of a STREAMS message. A module maximum packet size of **INFPSZ** (defined in **/usr/include/sys/stream.h**) defaults the maximum packet size to this value. The default is 0.

STRMCTLSZ

The maximum number of bytes allowed in the control portion of a STREAMS message. The control part of a message created with **putmsg** is not subject to the constraints of the minimum or maximum packet size, so this value is the only way of providing a limit for the control part of a message. The default is **1024**.

NMUXLINK

The maximum number of active multiplexors that may exist at any one time on the system. The default is **1024**.

NLOG
The maximum number of log devices available. The default is **16**.

NPIPE
The maximum number of streams pipe devices available. The default is **64**.

BSIZE
The maximum number of log messages allowed to be queued on a log driver's read queue. The default is **20**. This prevents the queue from being swamped by log messages during peaks.

NSTREVENT
The maximum number signal delivery requests (established with the **I_SETSIG** ioctl command) allowed to be queued on the signal event list. The default is **2048**.

Semaphore Configuration Variables

These semaphore configuration variables are also listed in **/usr/etc/master.d/dgux**. The variables are:

SEMMNI
The maximum number of unique semaphore sets that may be active at any one time on the system. The default is **1024**.

SEMMSL
The maximum number of semaphores that a semaphore set may contain. The default is **256**.

SEMOPM
The maximum number of semaphore operations that can be executed per **semop**(2) system call. The default is **10**.

SEMVMX
The maximum value a semaphore may have. The default is the maximum value for this parameter, **32767**.

SEMUME
The maximum number of undo entries per undo structure. The default is **10**.

SEMAEM
The maximum value of the adjustment for adjust-on-exit. The value is used whenever a semaphore value becomes greater than or equal to the absolute value of **semop**(2), unless the program has set its own value. The default value is the maximum value for this parameter, **16384**.

SEMAPM
The maximum number of processes that may specify adjust-on-exit at one time. The default is **16384**.

Shared Memory Configuration Variables

The tunable parameters shown below are associated with interprocess communication shared memory. These parameters are also defined in the **/usr/etc/master.d/dgux** file.

SHMMNI
The maximum number of shared memory identifiers system wide. Each entry contains 52 bytes. The default is **1024**.

SHMSEG
The number of attached shared memory segments per process. The default is **256**.

SHMMAX
The maximum shared memory segment size in bytes. The default is **(4*1024*1024)**.

SHMMIN
The minimum shared memory segment size. The default is **1**.

Message Configuration Variables

These parameter variables are also listed in **/usr/etc/master.d/dgux**. They set the message parameters shown in the following list.

MSGMNI
The maximum number of message queues that may exist in the system at one time. The default is **1024**.

MSGTQL
The maximum number of outstanding messages that may exist in the system at one time. The default is **1024**.

MSGMNB
The maximum number of bytes that a message queue may contain. The default is **4096**.

MSGMAX
The maximum number of bytes that a message may contain. The default **2048**.

Managing User Licenses

This section explains DG/UX system user licenses, how to get information about a user license, and how to upgrade the user count for a user license.

The user license defines the approved number of users that can run the DG/UX system concurrently on an AViiON server or workstation. A daemon monitors the number of concurrent users. When the number of concurrent users goes above the approved number, the daemon makes an entry in the **/var/license/usermon.log** log file and the user **root** receives email containing instructions on contacting Data General to increase the approved user count. The email message is sent to **root** once every hour as long as the user count exceeds the licensed number of users.

The following sections explain how to use **sysadm** to get information about the user license and how to upgrade a user license after Data General approves an increase in the user count for the license.

Listing User License Information

Use the **sysadm** operation `System-> License-> List` to view license information. You select one of the following options under List to specify what you want to see:

`Upgrade Info`

Lists the information you'll need before you contact Data General and who to contact to request a license upgrade. Much of the information is the same as you provide on a standard order form, such as your company name and an address for invoicing. One item is directly related to the user license:

`Current licensed user count`

The number of users currently approved to run the DG/UX system concurrently on the server or workstation.

Make sure that you have all of this information before you contact Data General to request an upgrade of the user license.

`Current License`

Shows the number of concurrent users approved for the current DG/UX system license and the number of current users:

```
Number of licensed users = 16
Number of licenses checked out = 15
```

`Current Users`

Lists the users recognized by the user count daemon as currently running DG/UX and shows information about these users:

```
User      pid     TTY      Hostname
jones     16322   pts/5    joe
pressley  9248    pts/4    fireball
palmer    2850    tty04
knightp   1145    pts/2    liberty
root      20184   tty14
```

Upgrading a User License

Use the **sysadm** operation `System-> License-> Upgrade` to upgrade a user license after Data General approves an increase in the user count. At the prompts, enter the license token given to you by Data General after approval of your upgrade.

When the Upgrade operation finishes, the new user count for the user license is in effect immediately.

End of Chapter

Chapter 5
Managing Software

The software packages on your system fall into the following categories:

- Operating system releases for OS clients, which are loaded into release areas.
- Software packages that conform to DG/UX system load and setup guidelines.
- Software packages that conform to 88open Consortium load and setup guidelines.
- Software packages that provide their own customized **sysadm** menus and operations.

The **sysadm** utility provides a menu and operations for each category.

Managing Release Areas

A release is a directory environment, or *release area*, that contains software that provides operating system service for OS clients. We distinguish between the *primary* release and *secondary* releases. The primary release runs as the operating system on OS servers and stand-alone systems. Any other releases are secondary releases.

For each release installed on your system, there is a directory in the **/srv/release** directory structure. For the primary release, there is **/srv/release/PRIMARY**, and for each secondary release, **/srv/release/***release_name*. The purpose of each release area is primarily to hold the root structures for any OS clients using the release and to hold one common copy of the **/usr** structure.

The primary release area is an exception because it does not hold the OS server system's root and **/usr** file systems. Instead, it contains links pointing to the system's root and **/usr** file systems.

The **/usr** file system contains host-independent programs and data files that users typically do not change. The rationale for this organization is to put in one place all of the operating system components that do not vary among systems using the same release of DG/UX; consequently, an OS server and any of its OS clients attached to the primary release may save disk space by sharing the same **/usr** file system.

A system's root file system, on the other hand, is the directory that contains data files, configuration files, and programs (such as kernels) that may vary from system to system; therefore, each system needs to have its own root file system. On OS servers and stand-alone systems, the root is the / directory. OS clients, on the other hand, have root directories created for them under the **/srv/release/***release_name***/root** directory. OS client root directories are based on a prototype found in **/usr/root.proto**.

The sysadm menu `Software -> Release Area` contains the operations you use to create, delete, and list release areas. The following sections elaborate on these tasks.

Creating a Software Release Area

To support OS clients, first you need to create the release area that will hold the OS client software. You do not need to create a release area for OS clients that will use the primary release because the primary release area already exists, as **/srv/release/PRIMARY**. See Chapter 6 for information on adding OS clients to the primary release. This chapter covers adding a secondary release.

A release is a collection of software packages intended for a specific architecture and operating system. Adding a release means creating the appropriate directories and files that will be used by the release. Once you have added a release area, you can load software into it.

Before beginning this procedure, make sure you have enough disk space in the file system containing the **/srv/release** directory structure. For the software's space requirements, consult the release notice for the software package. Use the **df**(1M) command to display the free space in a file system. Remember that the file system reserves a 10% free space buffer. If the file system is not big enough, you need to allocate more disk space by creating one or more additional virtual disks, creating file systems on them, and mounting them in some appropriate place under **/srv/release**.

For example, if you were going to add a release called **dgux5.4R3.00**, you would need to make sure that **/srv/release/dgux5.4R3.00** had enough space for:

- one copy of the DG/UX system's **/usr** file system and
- the OS clients' root file systems.

You could create and mount a file system at **/srv/release/dgux5.4R3.00/usr** and another at **/srv/release/dgux5.4R3.00/root**.

Chapter 7 explains how to create virtual disks. Chapter 9 explains how to create, add, and mount a file system.

Use the **sysadm** operation `Software -> Release Area -> Create` to create a release area, but before you invoke it, be prepared to answer these questions:

- What will you call the release area?
- Where will you put the root directories of the OS clients of this release?
- Will you establish a directory containing software (other than **/usr**-type OS software) that all the OS clients will share? If so, where will you put this directory?
- Where will you put the OS clients' swap areas?
- Where will you put the OS clients' dump areas?

Based on your answers, the system creates the following directories and file:

/srv/release/*release_name***/usr**
: For executables and data files that individual OS clients will not need to change.

/srv/release/*release_name***/usr/root.proto**
: The prototype for the root directories of OS clients using this release. The operation for adding an OS client makes a copy of the prototype root for each new OS client. OS clients are free to customize their own root directory.

/srv/share
: A directory that you can use at your own discretion, intended to contain whatever programs or files your OS clients may hold in common.

/srv/release/*release_name***/root**
: The directory containing the root directories of OS clients. The operation for adding OS clients creates a root directory here for each new OS client, naming the directory after the OS client's host name.

/srv/swap
: The directory containing the swap areas of OS clients. OS client swap areas are named after the host name of the OS client.

/srv/admin/releases/*release_name*
: The file containing a list of directories associated with this release.

/srv/dump/*client_name*
: The directory containing the dump files for OS clients. OS client dump files are named after the host name of the OS client.

At this point, you have created an empty release area. Using the operations under the Package menu, you now need to load the software for the release and set it up. After you have installed the software, add the OS clients. See Chapter 6 to add OS clients.

Deleting a Release Area

To delete a release area, use the **sysadm** operation `Software -> Release Area -> Delete`.

Deleting a release means deleting the release directory tree and removing files used by **sysadm** for the given release. You can only delete releases that no OS client is using. You cannot delete the PRIMARY release with this function.

Deleting a release removes the following directories and file from your system:

- **/srv/release/***release_name***/usr**
- **/srv/release/***release_name***/root**
- **/srv/admin/releases/***release_name*

Listing Release Information

To display the name of a release area and the pathname of the file system containing its host-independent (that is, **/usr**-equivalent) software, select the sysadm operation

`Software -> Release Area -> List` operation. You can display information on all releases or on a specific release.

Handling Packages Conforming to DG/UX Standards

The Package menu provides operations for handling software packages shipped by Data General for installation on DG/UX systems. The Data General ONC™/NFS® network software package and the X Window System software package are examples of Data General packages that you might load with these operations. Packages from other sources may also conform to the DG/UX package guidelines, which are discussed in *Porting and Developing Applications on the DG/UX™ System*.

Before you load any software package, you should consult the package's release notice. The release notice should contain vital information such as prerequisite software packages, where the package will load, and how much disk space it will require.

Typically, optional DG/UX packages load into a directory in **/usr/opt**. This directory should actually be the mount point for a virtual disk and file system created exclusively for holding the package. The reason for creating a dedicated file system like this is to avoid over-filling the **/usr** file system and to make future upgrades of the DG/UX system (which may involve overwriting **/usr**) less complicated.

All packages require that you load them, and many also require that you set them up. Loading involves simply transferring the software files from the distribution media to disk. Setting up a package involves running scripts provided with the package. These scripts may execute without requiring user interaction, or they may require that you supply some information necessary to initialize the software. Loading and setting up are both generally quite easy because the typical DG/UX package contains scripts to automate the procedures.

The Package menu provides operations for installing, loading, setting up, and listing DG/UX packages.

Installing Software into a Release Area

Install software packages with the **sysadm** operation `Software -> Package -> Install`.

This operation is intended for packages that conform to the DG/UX package guidelines as specified in *Porting and Developing Applications on the DG/UX™ System*.

Installation involves loading and setting up the software. Loading the software means transferring it from the media, such as tape, to disk. Setting up the software means running the setup script provided with the package to perform tasks such as initializing databases or querying you for needed information, for example. The

sysadm utility handles loading, and scripts provided with the software may handle the setup procedures.

NOTE: During the setup phase of package installation, the Install operation executes only the setup script having the same name as the package, if any. For example, if you install a package named Officeware that includes setup scripts named Officeware, Spreadsheet, and Wordprocess, the Install operation will run only the Officeware setup script. The operation does not run any other setup scripts; rather, at the end of the operation, it notifies you that other scripts exist. To execute them, select the Setup operation.

Before installing any software package, consult the package's release notice. The release notice contains vital information about prerequisite software, the location on disk where the software will reside, the amount of disk space that the package requires, and so forth. The release notice should tell how to install the package. The instructions in this manual are general and may not apply completely to the package that you are loading.

Before installing a software package, you need to know into which release area to load it. A release area is a directory structure intended to contain operating system software to support OS clients. Every system has a primary release area, which contains the currently installed release of the DG/UX system. If you are installing the package for use by OS clients that do not use the primary release, see "Installing for OS Clients" later in this chapter.

A Caution about Disk Space

Be careful not to load a package into a file system that does not have enough free space. To verify how much space a file system has, use the **df**(1M) command. The following example shows how to list the amount of free space in your **/usr** file system.

```
% df -t /usr ↵
/usr  (/dev/dsk/usr     ):   22610 blocks   27253 files
                   total:  240000 blocks   34560 files
```

The **df** output above shows that the **/usr** file system has 22,610 512-byte blocks of free space, or approximately 11.5 Mbytes.

If the file system does not have enough space, expand the file system with the **sysadm** operation `File System -> Local Filesys -> Expand`. Remember that you cannot boot from a file system built on an aggregated virtual disk that spans multiple physical disks; therefore, do not expand the / (root) or **/usr** file systems if doing so would add another partition on a separate physical disk. If you need to add space to / or **/usr**, first make sure that the desired amount of free space exists on the same physical disk.

Remember that for whatever amount of physical disk space you allocate, approximately 7% will be used for file system overhead, and 10% will be the free space buffer. Thus, you should allocate 17% more disk space than you intend to use. For example, if you need to allocate space for a 50 Mbyte software package, allocate 59 Mbytes of disk space.

See Chapter 9 for information on expanding file systems.

If you attempt to load the package into a file system that does not have enough space, the operation will fail. Depending on which file system it is that becomes full, you may also disrupt any current users' access to the file system. The disruption can be particularly inconvenient if you fill up the / or **/usr** file systems. If the operation fails, it will leave the software package partially loaded on disk, and you will have to remove the loaded files yourself. To find out what files to remove, see the release notice that shipped with the package.

Installing for OS Clients

If you are installing software for OS clients (diskless workstations) that use an operating system release other than the primary release, you will have to supply the appropriate release name when you install the software. The release area must already exist. If your OS clients use the same release of the DG/UX system that your OS server uses, you may install the package in the primary release area.

The installation operation loads any / file system files not only into your system's / file system but also into the OS client prototype / file system (which appears on the OS server as **/usr/root.proto**). Loading a software package on your system provides the option of loading it into the prototype / file system as well.

The prototype / file system is what **sysadm** uses as the model / file system when creating a new OS client root area. The operation you use to add an OS client copies the prototype root when making the initial root file system for the new OS client. This way, the OS client gets a copy of the operating system software as well as any other loaded packages.

Loading Software into a Release Area

To load software from the media to disk without setting it up, use the **sysadm** operation `Software -> Package -> Load`.

For the specified release, this operation loads the software into the / and **/usr** file systems. The operation also loads the root portions of the package into the prototype root and any existing OS client roots so that systems with OS clients will also have the package.

Read the software release notice before loading to see where the software will load and how much disk space it requires. Make sure your system has enough free disk space in the appropriate areas. If you need to create additional file systems for the package, make sure they are mounted before loading the package.

If you create new file systems for a software package, OS clients who desire access to the package will have to mount the new file systems onto the appropriate directory on their system.

Setting Up Software in a Release Area

To set up software that you have already loaded, use the **sysadm** operation `Software -> Package -> Set up`.

Setting up software may consist of any number of steps, all determined by the setup script, if any, loaded with the package. The setting up process involves initializing any files as necessary and distributing files to their required locations on the system. The setup script often prompts you for information necessary to customize the software for your site. There may also be other steps you need to complete before the software package is usable; consult the package's release notice for more information.

For a given release area, this operation sets up the software in both the **/usr** and root file systems. The operation also lets you set up the software in the roots of OS clients attached to the release. Whether or not you set the software up in OS clients' roots depends on the nature of the given package's setup script and whether or not the users of the OS clients prefer to set up their own systems. This operation locates all setup scripts that have not been run from a software package and allows the user to execute them.

Listing Packages

To display information about software packages on release media or installed on your system, select the **sysadm** operation `Software -> Package -> List`.

When listing packages installed on a system that has multiple release areas, you need to specify which release area's packages to list. You may choose to list any or all of the packages that have been installed in that release area.

The default listing shows only the short name of each loaded package. For packages loaded on the system, the detailed listing shows the short name, the complete name, the release number of the package, and the date the package was created or released. For packages on release media, you see the table of contents, which includes the package's **tar** or **tarZ** components, component sizes, component load points, and so on.

Handling Packages Conforming to 88open Consortium Standards

In addition to installing packages that conform to the DG/UX software package guidelines, you may install packages that conform to the 88open Consortium standard for software packages. Packages that conform to the 88open standard are identical in structure to AT&T System V Release 4 AIS packages.

The 88open Package menu provides operations for adding, deleting, and listing information on packages that comply with the 88open standard.

Adding 88open Packages

To install packages that conform to the 88open Consortium standard for software package installation, use the **sysadm** operation `Software -> 88open Package -> Add`.

See the package's release notice for vital information such as where the package loads, how much disk space it requires, and any other procedures required to make the package usable.

This operation loads the software into the appropriate system or spool area (a directory for holding software that is loaded but not installed). If you install the package on the system, the package is immediately available for use, provided you have completed any other installation or setup procedures described in the package's release notice. Installing a package in the spool area does not allow you to use the package immediately. To make the package usable on your system, you need to invoke the Add operation again, this time specify the spool area as the device that contains the package.

Loading a package into the spool area has its advantages because it provides you with a convenient means of:

- Storing the package on-line until you are ready to complete installation on the system.
- Copying the package by allowing you to load the package from the spool area back to another portable medium such as a tape.
- Making the package available to other systems (such as OS clients) who can mount the spool directory and install the package themselves.

To remove an installed package, use the Delete operation.

Deleting 88open Packages

To remove an installed 88open package, use the **sysadm** operation `Software -> 88open Package -> Delete`.

The operation removes files associated with the package and reverses any changes that the package made to other system files during installation.

The exact behavior of the Delete operation depends on scripts and file listings installed with the package that you are removing; therefore, the behavior of the Delete operation may vary from package to package. If you have created any of your own files or directories in the installed package's directories, you may want to move them (if you intend to save them) before deleting the package.

You may delete packages either from the system or from the spool area.

Listing 88open Packages

To display information about installed 88open packages, use the **sysadm** operation `Software -> 88open Package -> List`.

This operation lists 88open packages on devices listed in **/etc/device.tab** or loaded on your system. By default, **/etc/device.tab** includes an entry for a tape device at **/dev/rmt/0** and an entry for the packaging spool area, **/var/spool/pkg**.

When you list packages, you may select either a detailed listing or a simple listing.

Handling Other Applications

The Applications Management menu provides menus and operations for any additional software that you have installed on your system beyond just the DG/UX system and bundled software. Other software packages may or may not add menus and operations to the Applications Management menu. For more information, see the documentation for your other software packages.

End of Chapter

Chapter 6
Managing OS and X Terminal Clients

This chapter tells how to manage OS (operating system) and X terminal clients. An OS client is a workstation that depends on another system, called the OS server, for its operating system software. The OS client boots its network controller, thereby initiating a series of network transactions with its OS server. The OS server transfers a kernel image to the OS client so the client can boot and mount remote file systems. The OS client operations apply only if you are the administrator of an OS server system.

An X terminal client is a type of graphics terminal that provides X Window System graphics support even though it does not have its own CPU. Lacking its own disk and operating system, the X terminal depends on a server in somewhat the same way that an OS client does.

The **sysadm** Client menu provides menus containing operations for managing OS and X terminal clients. Of the two main sections of this chapter, the first section covers OS clients, and the second section covers X terminal clients.

Managing OS Clients

The OS Client menu provides operations for adding a client to a release, deleting a client from a release, modifying a client's bootstrap link, listing clients attached to a release, and setting a client's current release. The OS Client menu also provides the Defaults menu for managing the values that appear as the defaults when you add a client.

Adding an OS Client to a Release

You should create a kernel for an OS client before adding the OS client. Use the **sysadm** operation `System -> Kernel -> Build`, which is discussed in Chapter 4.

To add an OS client, select the Add operation. You may attach a client to multiple releases by performing the Add operation once for each release. Use the Set operation, discussed later, to select which release the OS client boots.

After a release has been set up on a server (see the **sysadm** Software menu), you can add an OS client to that release. Adding an OS client means creating a root directory for the client and then recording information about the client. Here are the steps for adding an OS client:

1. If you wish to define a set of default values for the queries in the Add operation, use the `Client -> OS Client -> Defaults -> Create` operation. Use

`Client -> OS Client -> Defaults -> Set` operation to make the defaults set the current set.

2. Add the client's Internet address to the **hosts** database with the operation `Networking -> TCP/IP -> Databases Hosts -> Add`. Get the Internet address from your network administrator. Then add the client's Ethernet address to the **ethers** database with the operation `Networking -> TCP/IP -> Databases -> Ethers -> Add`.

   Some AViiON computers display their Ethernet address when you turn on power. See the Networking menu for more information on the **sysadm** operations for adding network database entries. For more information on the networking packages, see *Managing TCP/IP on the DG/UX™ System* and *Managing ONC™/NFS® and Its Facilities on the DG/UX™ System.*

3. Make sure the following directories have sufficient space as noted:

   /srv/swap

   This directory is for the files that OS clients use as their swap areas. The file system containing this directory must be large enough to hold the swap areas for *all* OS clients that your system serves.

   Use the following formula to calculate OS client swap space. The number of blocks–per–OS–client is either 50,000 (24 Mbyrtes) or 1.5 times physical memory, whichever is larger.

   ```
   ( number-of-OS-clients * blocks-per-OS-client ) + 7% overhead

   ( 4 * 50,000 ) + 7% = 200,000 + 7% = 214,000 blocks
   ```

   If your swap space is not sufficient once your system is running, use the free swap value reported by the **sar** command to determine the available swap space. (Divide the amount of free swap space by 2048 to convert blocks to megabytes.) You should maintain the amount of free swap space at 15% to 30% of the total physical memory plus swap area space on the system. Refer to *Customizing the DG/UX™ System* for more information on calculating OS client swap space.

   /srv/dump

   This directory is for memory dump images that OS clients may produce in the event of a system panic or any time the client administrator chooses to make a memory dump. The file system containing this directory should be large enough to hold at least one system memory image of an OS client. A system's memory dump is the same size as its physical memory, or, if taking only a kernel dump rather than a complete dump, about half that size. This directory is optional but recommended.

   /srv/release/*release***/root**

   This directory is for the root directories of the OS clients. You need one such directory for each release supported by your OS server. A DG/UX system root directory requires 40,000 512-byte blocks. The file system containing this directory must be large enough to contain the root directories for all OS clients attached to the release.

/srv/release/*release***/usr**

This directory is for the host-independent (**/usr**-equivalent) directory of a secondary release. This directory is intended to contain the portion of an operating system that does not vary from system to system within that release. You do not need this directory if all of your OS clients use the PRIMARY release. You need one such directory for each secondary release supported by your OS server. The file system containing this directory needs to be large enough to hold the host-independent software for the release.

If the file systems that will contain the above directories are not large enough, you may need to increase the size of the file systems or create new file systems. See Chapter 7 for information on creating virtual disks; Chapter 9 for information on adding a file system to the **/etc/fstab** file and expanding the size of a file system.

4. Perform the Add operation for each client.
5. Set up software packages on the client system. See the Software menu.

NOTE: Because of basic operating system differences, the **sysadm** operations for adding DG/UX OS clients may not completely set up foreign OS clients. AViiON servers will support foreign OS clients, but Data General cannot supply the foreign system-specific information necessary for a complete setup. Consult the foreign system's documentation and/or your Data General representative.

The **sysadm** operation `Client -> OS Client -> Add` requires that you supply the following information:

`Client Host Name`

This is the name by which the OS client will be known on a network. You should have already added **hosts** and **ethers** entries for the client in your TCP/IP databases.

`Release Area`

This is the operating system release for the OS client. You should have already added this release using operations in the Software menu. Adding a client to a release does not set the client to boot that release if the client is already set to boot a different release. You set a client's current release with the Set operation. The default release, **PRIMARY**, already exists because it is the currently installed release of the DG/UX system on the server. Make sure the file system containing the release has enough free space for the new root directory. A DG/UX system root requires 20 Mbytes (40,000 512-byte blocks) by default.

`Server Host Name`

This is the name of the system that will be the client's OS server. The reason you have a choice here is because systems with more than one network device have more than one host name. A system connected to two different networks, for instance, has one host name for one network interface and another host name for the second network interface. You should select the host name for the network interface connected to the client's network.

`Home Directory`

This directory is the one that will contain the home directories of users on the client. On the system where the directory resides, the directory must be exported. The Add operation modifies the client's file system table (**/etc/fstab**) to make the directory available on the client system.

`Swap Size`

This is the amount of swap space that you want to allocate for the client system. There is no formula for calculating the ideal amount of swap space for a system. A general rule is to allocate 1.5 times the amount of physical memory in the client computer. The default physical memory size, 16 Mbytes, is about 32,000 512-byte blocks. Multiplying 1.5 by 32,000 blocks produces 48,000 blocks, which rounds up to 50,000 blocks. Make sure the file system containing the **/srv/swap** directory has enough free space to hold the swap file.

`Kernel Pathname`

This is the pathname of the kernel image that the client will use. OS client kernels generally reside in **/srv/release/***release_name*/**root/_Kernels**.

`Bootstrap File`

This is the pathname of the bootstrap file that the client will use. The default, **/usr/stand/boot.aviion**, is for the PRIMARY release.

Files Created by the Add Operation

The Add operation copies or modifies a number of files. This section outlines some of these changes.

When you add a client, the client inherits the server's environment. This means that the client receives information from the server's various parameter files: **dgux.params**, **tcpip.params**, **nfs.params**, and so on. Clients can modify these as they wish. The Add operation creates a number of directories and files for a new client: a root directory, a swap file, client parameter and data files, a link to a common OS client kernel, a link to a common OS client secondary bootstrap, an entry in the server's **bootparams** file, entries in the server's **exports** file, and client/release data files used by **sysadm**.

Client Root

An OS client's root space is created on the server. For the **PRIMARY** release, **sysadm** copies the files in **/srv/release/PRIMARY/usr/root.proto** to the client root area. The client root area is **/srv/release/***release_name*/**root/***client*, where *release_name* is the name of the release and *client* is the client's host name.

Make sure the file system containing **/srv/release/***release_name*/**root/***client* has enough free space to hold the client root. The root directory for a DG/UX system requires 20 Mbytes (40,000 512-byte blocks) by default.

Client Swap File

The client swap file is **/srv/swap/***client*. Make sure the file system containing **/srv/swap** has enough free space to hold the swap file for all OS clients. The swap file is fixed at the size you define.

Client fstab File

The Add operation adds the following entries to the client's **/etc/fstab** file, whose pathname is **/srv/release/***release_name***/root/***client***/etc/fstab** on the server:

```
server:/srv/release/release_name/root/client /         nfs  rw x 0
server:/usr                                  /usr      nfs  ro x 0
server:/srv/swap/client                      swap      swap sw x 0
server:home_dir                              home_dir  nfs  rw x 0
```

where *server* is the server's host name, *client* is the client's host name, and *home_dir* is the file system containing the home directory of the client's users.

The Add operation adds the entry for the client's home directory only if the home directory is on the server. **fstab** should contain entries for all other file systems that a client needs to access. See **fstab**(4).

Client nfs.params File

The Add operation modifies the client's **/etc/nfs.params** file, whose pathname is **/srv/release/***release_name***/root/***client***/etc/nfs.params** on the server, and sets the NIS (Network Information System) domain name to be the same as the server's NIS domain name. For more information on NIS, see *Managing ONC™/NFS® and Its Facilities on the DG/UX™ System.*

Client Kernel Link

An OS client boots **/srv/release/***release_name***/root/_Kernels/dgux.diskless** by default. The Add operation links this file, if it exists, to a file in the client's root directory, **dgux**. This kernel does not exist on the DG/UX system as shipped. If you intend for your OS clients to use this kernel, you need to build it. Use the operation `System -> Kernel -> Build`, discussed in Chapter 4.

Server Bootstrap Link

An OS client uses a secondary bootstrap to load **dgux.diskless** over the network. This default bootstrap file is **/usr/stand/boot.aviion**. The Add operation makes a symbolic link from **/usr/stand/boot.aviion** to **/tftpboot/***client_ip_addr*, where *client_ip_addr* is the client's Internet address in hexadecimal. The Add operation gets the client's Internet address from the **/etc/hosts** file.

Server bootparams File

The Add operation puts the following entry in **/etc/bootparams** for an OS client on an OS server:

```
client  root=server:/srv/release/PRIMARY/root/client \
        swap=server:/srv/swap/client \
        dump=server:/srv/dump/client
```

where *client* is the client host name, and *server* is the server host name.

The entry is actually a single logical line. In the entry above, backslashes at the ends of the first two lines escape out the New Line characters, effectively removing them and making the three physical lines one logical line.

Server exports File

The Add operation puts the following entries in **/etc/exports** for an OS client on an OS server:

```
/srv/release/PRIMARY/root/client  -access=client,root=client
/srv/swap/client  -access=client,root=client
/srv/dump/client  -access=client,root=client
```

where *client* is the client host name.

If the **/usr** file system is not already exported, the operation adds it to the server's **/etc/exports** file.

Client/Release Data Files for sysadm

The Add operation creates **/srv/admin/clients** and **/srv/admin/releases**. These directories contain files used by the **sysadm** program. Do not modify the contents of these directories or files yourself.

Deleting an OS Client from a Release

Select the **sysadm** operation `Client -> OS Client -> Delete` to remove a client from a release.

Deleting a client means deleting a client's root directory tree for a given release. Also, **sysadm** information on a client/release pair is erased. This operation displays each client/release pair and asks if it should be deleted.

The Delete operation prompts you for the name of the release. Optionally, you may choose to save the client's root file system.

Modifying an OS Client's Bootstrap Link

Select the `Client -> OS Client -> Modify` operation to change the pathname of the file used as a client's secondary bootstrap. The Modify operation prompts you for the client name, the release name, and for the new bootstrap file pathname.

Listing an OS Client Information

You may use the operation `Client -> OS Client -> List` in two ways. You can display general information about all clients by specifying **all** (the default), or you can display detailed information about a single client.

Changing an OS Client's Boot Release

In the case where a client is attached to more than one release, the `Client -> OS Client -> Set` operation allows you to change the default boot path. Adding a client to a new release with the `Client -> OS Client -> Add` operation does not change the client's default boot path; you must change the default boot path explicitly with the Set operation.

Managing OS Client Defaults Sets

The **sysadm** operation `Client -> OS Clients -> Defaults` provides operations for managing sets of default values that appear when you use the Add operation to add a client to a release. The menu provides operations to create, remove, modify, and list defaults sets. There is also an operation to select which defaults set is currently used by the Add operation.

Creating a Defaults Set

When you add an OS client, **sysadm** requires that you set several parameters. As it often does, **sysadm** offers default values for the parameters. If the defaults are not what you want, however, you may supply instead your own defaults in what is called a defaults set. Use the **sysadm** operation `Client -> OS Client -> Default -> Create` to create such a defaults set.

Once you have created the defaults set, use the `Modify` operation to set the values for the defaults set. Use the `Set` operation (in the OS Client operation) to set the current defaults set.

Removing a Defaults Set

Use the **sysadm** operation `Client -> OS Client -> Defaults -> Remove` to delete an OS client defaults set created with the Create operation. Deleting a defaults set has no effect on clients added using the default set.

Modifying a Defaults Set

Use the **sysadm** operation `Client -> OS Client -> Defaults -> Modify` to set or change the values in an OS client defaults set created with the Create operation. A defaults set includes the same parameters that you see in the Add operation of the OS Clients menu.

The parameters you need to specify to modify a defaults set are:

- Set Name
- Release Area
- Server Host Name
- Home Directory
- Swap Size
- Kernel Pathname
- Bootstrap File

See the section on the OS Client menu's Add operation for a discussion of each parameter. Modifying a defaults set has no effect on OS clients added when the defaults set was in use.

Listing Defaults Sets

Use the `List` operation to display the parameters for an OS client default set. See the section on the OS Client menu's Add operation for a discussion of each parameter.

Selecting a Defaults Set

Use the `Select` operation to set the current OS clients defaults set to be used by the Add operation. Adding an OS client defaults set does not automatically make that set the current one; you need to use the Select operation to make it current.

Managing X Terminal Clients

In addition to OS clients, which are typically diskless workstations, the DG/UX system supports the AViiON AVX-30 network display station, referred to as an X terminal. An X terminal is a terminal with X Window System graphics capabilities. Unlike a normal user terminal, which connects to a computer's asynchronous communication port, the X terminal connects directly to the network.

The X terminal boots in a manner that is similar to an OS client in that it starts by announcing its need to boot over the network. The server assigned to the X terminal then transfers bootstrap code to the X terminal, allowing the X terminal to initialize itself and become an active member of the network. Unlike an OS client, an X terminal does not run an operating system that it receives from the server; instead, it runs its own network and graphics software sufficient to communicate with other hosts on the network and provide an X Window System graphics environment for the user.

The X Terminal menu provides operations for adding, deleting, modifying, and listing profiles for X terminals supported by your system. The Defaults menu is for managing different values for the default bootstrap that appears when you add an X terminal client.

Adding an X Terminal Client

Select the **sysadm** `Client -> X Terminal -> Add` operation to add support for an X terminal client. This operation sets your server up so that an AViiON AVX-30 network display station (or similar X terminal) can boot. This operation sets up certain files so that when the X terminal boots, your server can send it the proper bootstrap file to get it started. For more information on the AVX-30 network display station, see the AVX-30 network display station documentation and software release notice.

Before performing the Add operation, add the X terminal client's Internet and Ethernet addresses to your TCP/IP databases. See the Networking menu. The X terminal displays its Ethernet address when you turn on power. Get the Internet address from your network administrator.

The Add operation prompts you for the host name by which the X terminal client is known on the network. The operation also prompts you for the pathname of the

bootstrap file needed to boot the X terminal. The default bootstrap file is intended for AViiON AVX-30 network display stations. For other X terminals, you need to specify the pathname of a bootstrap file that is appropriate for that X terminal.

Deleting an X Terminal Client

Select the **sysadm** operation `Client -> X Terminal -> Delete` to delete the files on your server that allow an X terminal client (such as the Data General AVX-30 network display station) to boot. Select **all** to delete all X terminal clients, or specify the name of a single X terminal client.

After deleting an X terminal client, you may wish to remove the client's Internet and Ethernet addresses from your TCP/IP databases. See the Networking menu.

Modifying X Terminal Clients

Select the **sysadm** operation `Client -> X Terminal -> Modify` to modify the Modify operation to change the bootstrap file pathname for X terminal clients supported by your system. To change the bootstrap file for all clients, select **all**. To change the bootstrap file for just one X terminal client, select the client's host name.

Listing X Terminal Clients

Select the **sysadm** operation `Client -> X Terminal -> List` to display parameters associated with the X terminal clients supported by your server. The fields in the display are:

`Client Name` The host name by which the X terminal client is known on the network.

`Address` The client's Internet address.

`Bootstrap` The client's Internet address, in hexadecimal, which is used as a link file pointing to the bootstrap file.

`Linked to` The pathname of the bootstrap file to which the link points.

Setting X Terminal Client Defaults

The `Client -> X Terminal -> Defaults` menu provides operations for creating, removing, modifying, and listing default sets. A default set is a value (for the bootstrap file pathname) that can appear as the default when you add an X terminal client. An additional operation allows you to select which set determines the default currently used by the Add operation.

Creating a Default Set

Select the `Create` operation to make a new default set. Use the Modify operation to define the bootstrap value in the default set. Creating an X terminal client default set does not automatically make the set the current one; you need to use the Select operation to make it current.

Removing a Default Set

Select the `Remove` operation to delete an existing default set. Deleting a default set has no effect on X terminal clients added when the default set was in use.

Modifying a Default Set

Select the `Modify` operation to set or change the value of the bootstrap parameter in an existing default set. Modifying a default set has no effect on X terminal clients added when the default set was in use.

Listing Default Sets

Select the `List` operation to display the bootstrap parameter setting in an existing default set.

Selecting a Default Set

Use the `Select` operation to set the current X terminal clients default set to be used by the Add operation. Creating an X terminal client default set does not automatically make that set the current one; you need to use the `Select` operation to make it current.

End of Chapter

Chapter 7
Managing Disks

This chapter covers the tasks involved in managing your disks. These tasks include:

- Formatting physical disks
- Configuring and deconfiguring, registering and deregistering, copying, tracking bad blocks, and converting physical disks between logical and virtual disk format
- Creating and maintaining striped and normal (nonstriped) virtual disks
- Creating and maintaining mirrored virtual disks
- Creating and maintaining cached disks
- Setting up and maintaining failover disks
- Setting up disk arrays

Virtual Disk Management (VDM)

DG/UX 5.4R3.00 introduces the new virtual disk management (VDM) technology, which enables system administrators to manipulate storage space with the data online and in use. This on–line convenience, in turn, enables you to dynamically create and rearrange disk resources for improved system performance and high availability of data.

The most important new feature that VDM offers is high availability. You can do almost all disk management operations on–line, providing users and applications uninterrupted access to data. Such operations include expanding file systems on–line, moving partitions from one physical disk to another on–line, and creating or removing software mirrors on–line. You do not have to shut down applications to perform these tasks, nor do you need to schedule them during times when there are no users on the system. VDM enables you to manage disks at your convenience—even in the middle of the work day.

Both the virtual disk management technology and its predecessor, logical disk management (LDM), enable you to subdivide physical disks so that you can store many types of information on different parts of a single physical disk.

NOTE: Comparisons between VDM and its predecessor, logical disk management (LDM), will be meaningful to customers who are familiar with releases prior to DG/UX 5.4R3.00. Customers who have no experience with previous releases may ignore all discussions in which VDM is compared with its predecessor.

With disk managers, you can subdivide the physical disk into discrete amounts of space, each of which can be used to store a different file system or database. Furthermore, both VDM and LDM allow you to combine amounts of storage from multiple physical disks to form storage areas that exceed the size of a single physical disk. They also enable you to replicate data through software mirroring and to set up fast caches.

VDM and LDM provide many of the same functions, but their implementations are significantly different. VDM can be thought of as the next generation of LDM, offering an expanded set of disk management options for you to use when organizing your data resources.

To create and maintain virtual disks, use the **sysadm** menu `Device -> Disk -> Virtual`. See the section on virtual disks in this chapter for more information.

NOTE: Be aware that DG/UX 5.4R3.00 supports VDM only, replacing entirely its predecessor, LDM. The current arrangement of **sysadm** menus and options is significantly different from releases prior to DG/UX 5.4R3.00.

Virtual Disk Terminology

While some configurations are common to both the logical and virtual disk technologies, there are some virtual disk configuration possibilities that simply don't exist for logical disks. For example, the lowest–level building block in each technology, a logical disk piece in LDM and a partitioned virtual disk in VDM, is roughly the same. However, VDM offers a hierarchical flexibility that allows you to vertically combine partitions with other partitions to form aggregations, which you can further combine with other partitions or aggregations to form more complex virtual disks. Because of the new VDM technology, new terms were needed. Their definitions and other relevant terms and definitions follow:

Virtual disk

VDM counterpart to the LDM logical disk, which is an amount of space that you reserve from a physical disk onto which data is loaded. A virtual disk is further characterized by a type: partition, aggregation, mirror, cache, or combination of same types.

Partition

A type of virtual disk. VDM equivalent of the logical disk piece, it is the specific, contiguous space reserved on a physical disk for later data storage. You can create a partition from a physical disk or partition from an existing virtual disk. The difference between a partition and a logical disk piece is that a partition can be mounted; a logical disk piece cannot.

Aggregation

A type of virtual disk. VDM equivalent of multi–pieced logical disk under LDM. The combination of virtual disks to form a "parent" virtual disk whose total size is the sum of the sizes of its "children" components. For example, you can form an aggregation by grouping multiple partitions, aggregations, mirrors or any combination of the same virtual disk types. An aggregation

can comprise up to 120 partitions on one or more physical disks. An aggregation can be striped as long as each child component is identically sized.

Mirror

A type of virtual disk. A collection of identically sized virtual disks that are duplicates (copies) of each other. Multi–image (three maximum) mirrors can be created from partitions, aggregations, caches, other mirrors, or any combination of the same virtual disk types. Mirroring maximizes data reliability and high availability.

Cache

A type of virtual disk. Association of a virtual disk on a typically slow, large storage device with a virtual disk on a fast device so that an application can use the fast device for read and write operations while the operating system duplicates these operations on the larger, slower one. Caching offers the benefit of speed from the fast device and capacity from the large device. A typical cache is formed by attaching a NVRAM card (fast device) or a fast disk to a slow one (disk device). Caching maximizes data throughput for applications that are read and write intensive.

Physical disk

A physical disk is the disk drive hardware. The SCM (the AViiON hardware's System Control Monitor) and the DG/UX operating system refer to physical disks or CLARiiON disk array units with DG/UX device names such as **sd(insc(),0,0)** and **cimd(1,1)**. For more information on device names, see *Customizing the DG/UX™ System*.

File system

A file system is a software–formatted portion of disk space that is typically located on a virtual disk. An exception is for diskettes, which use a file system that occupies the entire physical disk. The file system contains the internal data structures that the operating system requires to keep track of files and directories on the virtual disk. Typically, you build a file system on each virtual disk that you create except on virtual disks used as swap area (for demand paging) and any virtual disks to be used for direct access by applications such as database managers.

Volume

A virtual disk for which entries are created in **/dev/dsk** or **/dev/rdsk**. These entries are needed to mount or read and write the virtual disk. Any virtual disk can be declared a volume and therefore, be directly accessible. By contrast, logical disk pieces cannot be independently mounted.

Features of Disk Management

The DG/UX system offers a number of features that allow you to get better service from your disk file systems. These features can improve performance and data availability as well as provide compatibility with non-DG/UX file systems.

The following sections summarize these features.

Memory File Systems

For applications that would benefit from very fast access to relatively small databases, you can create memory ("ramdisk") file systems. A memory file system is a portion of your computer's physical memory, formatted as a DG/UX file system and mounted. You can access it the same as any file system. For more information on memory file systems, see "Creating a File System" in Chapter 9.

Mirrored Virtual Disks

You can improve the reliability and availability of your Data General AViiON system by mirroring virtual disks. Virtual disk mirroring involves maintaining redundant virtual disks where all are "mirror images" of each other; they all contain the same data. The system manages access to the disks in a manner that is transparent to users.

Disk mirrors provide higher data availability by allowing your system to continue service to users even when disk errors occur. Disk mirrors also protect data integrity by maintaining redundant images of the same data. In limited cases, mirrors can improve disk I/O throughput and thus improve overall system performance.

There are two ways to configure disk mirrors: through the hardware or through the operating system software. To configure hardware mirrors, you need a disk-array subsystem. This manual does not cover hardware disk mirrors. See the disk-array subsystem documentation for more information.

To create software disk mirrors, use the **sysadm** menu `Device -> Disk -> Virtual -> Mirrors`. See the section on mirroring in this chapter for more information.

Software Data Striping

Applications that perform a lot of random I/O (reads as well as writes) and applications that perform a lot of sequential reads can benefit from data striping. Data striping involves storing sequential file elements on alternating physical disks so that sequential disk accesses may be interleaved, taking advantage of the disk hardware's read-ahead feature to speed disk I/O.

From the administrator's point of view, a striped virtual disk is as easy as any other virtual disk to create and use because the system handles the data striping itself. Once you create a striped virtual disk, you cannot expand or shrink it. The section on creating virtual disks tells how to create software-striped disks and describes the kinds of applications that can benefit from data striping.

For information on hardware data striping, see the disk-array documentation.

Fast Recovery File Systems

To reduce the amount of time that the system will require to recover a file system after a failure, mount the file system with **fsck** logging turned on. With **fsck**

logging, the system logs file system modifications to reduce the amount of time that **fsck** requires to verify the integrity of the file system. This feature is desirable for systems where rapid recovery and high availability are crucial. For more information on fast recovery file systems, see the section on **fsck** in Chapter 9. This feature does not support software disk mirrors.

Write Verification

In applications where data integrity is vital, you can benefit from write verification. By turning write verification on for a physical disk, you can be sure that data written to the disk is readable.

You can enable write verification only for SCSI disks that support the feature. See your disk hardware documentation.

Normally, disk hardware does not verify that data just written to disk is readable. This behavior makes your data vulnerable to flaws in the storage medium because the disk has no way of detecting when it has just written to a flawed block on the disk, for example. With write verification turned on for a disk, however, the disk hardware verifies every write operation by reading the written data off of the disk and comparing it to the data as originally received. If the data read from disk does not match the data originally received in the write request, the hardware returns an error to the system. Thus, write verification ensures the integrity of your data.

The tradeoff with write verification is in performance. The additional verification overhead in the hardware can have a significant impact on the performance of write-intensive applications. You should experiment with your applications to see how write verification affects performance. Write verification has no effect on read operations.

To turn write verification on for a disk, use the **dkctl**(1M) command with the **wchk** option. First, you need to know the path for the physical disk drive, for example, **/dev/pdsk/0**. Look in the **/etc/devlinktab** file for a list of drives on your system. By comparing the disk's device name with the long names in **devlinktab**, you can find the correct entry and note the short name. Use the short name in the **dkctl** command line. See *Customizing the DG/UX™ System* for more information on device names.

For example, the following command line turns on write verification for disk **/dev/pdsk/0**:

```
# dkctl /dev/pdsk/0 wchk ↵
```

This command line enables write verification for the disk every time you boot the system. To enable write verification only until the next system boot, include the **-t** (**t**emporary) option:

```
# dkctl -t /dev/pdsk/0 wchk ↵
```

To turn write verification off for the disk, issue this command:

```
# dkctl /dev/pdsk/0 -wchk ↵
```

See the **dkctl**(1M) manual page for more information.

Cached Virtual Disks

Disk caching associates two virtual disks, typically one on a small, fast device (front end) with another on a large, slow device (back end), so that an application uses the fast device for read and write operations while the operating system duplicates these operations on the larger device. The purpose of the configuration is to accelerate file system access for I/O-intensive applications without risking data integrity.

See the section on managing cached disks later in this chapter for more information.

Device Sharing and Disk Failover

Disk failover and tape sharing involve setting up two systems in a *dual-ported* configuration or a *dual-initiator* configuration where they share a common SCSI bus or disk-array. Disk subsystems that support *dual-porting* can also provide failover. In these configurations, the two systems provide alternate paths to the same devices and data. The disk-failover feature provides a simple way of transferring control of a disk or disk-array from one system to the other. Disk failover thus provides not only a means of restoring database and application access quickly after a failure but also of balancing the I/O or CPU load shared by two normally functioning systems.

See the section on managing failover disks in Chapter 8 for more information.

CD-ROM, Diskette, and Magneto-optical Disk Drive Support

The DG/UX system supports a variety of SCSI disk drives including CD-ROM, diskette, and magneto-optical disk drives. For more information on these drives, see Chapter 15.

Non-DG/UX File System Support

In addition to the DG/UX file system, the DG/UX system supports MS-DOS, High Sierra, and ISO 9660 file systems. MS-DOS file system support is particularly useful for sites with diskette devices, while High Sierra and ISO 9660 support is useful in environments with CD-ROM disk drives. Chapter 15 covers diskette and CD-ROM drives and several other types of SCSI drives. Chapter 9 covers file system operations. These chapters also include instructions for configuring your kernel for non-DG/UX file systems and for creating and mounting MS-DOS file systems.

Support for Multiple VME Channels

The DG/UX system supports multiple VME channels on AV 9500 systems, which, when used, must be specified in the DG/UX common device specifications and the **vme()** name. This means you can specify explicitly a VME channel for a device; for example, **syac(vme(1),4)** refers to the fifth standard **syac** device attached to the second VME channel controller. All the following DG/UX VME devices allow (but do not require) the parent VME to be specified as their first parameter:

cied and **cimd**	**cien**	**cisc**	**dgsc**
hada	**hken**	**nvrd**	**pefn**
ssid	**syac**	**vitr**	**vsxb**

For compatibility with systems that do not use multiple VME boards, the **vme()** device name has been designed to be optional; if you omit it, the system will assume **"vme(0),"**. Therefore, the device name **sd(dgsc(vme(0),2),3,1)** is functionally the same as **sd(dgsc(2),3,1)**. When building an autoconfigured kernel, the system automatically generates entries in the system configuration file for VME channels (see Chapter 4 for information on kernel building). See *Customizing the DG/UX™ System* for more information on device naming.

Expanded SCSI Device Naming Requirements

You will need to use the fourth field (logical unit number) in the DG/UX common device specification when naming these types of devices:

- devices attached to a VME channel controller
- failover devices
- diskettes

The following syntax is used for these devices.

device **(** *controller–type* [*@device–code*] **,** [**(** [*vme*(*n*)] *controller–number* [**,***controller–SCSI–ID* **),** *device–SCSI–ID* [**,***LUN*] **)**

An example follows.

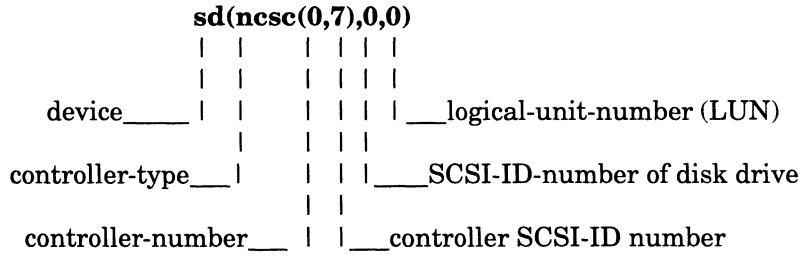

Refer to Appendix E for more details on using these fields. For complete information on device naming, refer to *Customizing the DG/UX™ System*.

Stand–alone and Stand–among sysadm

You perform disk management tasks using the **sysadm** operation `Device -> Disk`. For more information on **sysadm**, see the **sysadm**(1M) manual page or simply invoke the utility and select the Help menu. Alternatively, you may use the **admvdisk** and **admpdisk** commands to perform the **sysadm** operations, but from the shell. See the **admvdisk**(1M) and **admpdisk**(1M) manual pages for more information.

The **sysadm** utility appears on your system in two forms that have practically identical user interfaces: stand–alone **sysadm** and stand–among **sysadm**. Refer to Chapter 1 for the appearance of the stand–alone **sysadm** and stand–among **sysadm**

top–level menus. Although the interfaces for both are practically the same, there are four primary differences between the two:

- Stand–alone **sysadm** offers only an ASCII terminal interface while stand–among offers both an ASCII terminal interface and an OSF/Motif interface.
- Stand–alone **sysadm** offers a subset of the operations provided in stand–among **sysadm**. The stand–alone **sysadm** implementation is devoted to physical disk, virtual disk, file system, and software installation operations. The file system operations are slightly different between the two interfaces. In stand–among **sysadm**, you specify a file system by its mount point such as **/usr/opt/X11**, which is listed in the file **/etc/fstab**. Since the **/etc/fstab** file is unavailable when the system is not running, in stand–alone **sysadm** you express a file system's mount point using a different format **/mnt/***virtual–disk–name*. Thus, to check the **/usr/opt/X11** file system in stand–alone **sysadm**, you would specify **/mnt/usr_opt_X11**.
- Stand–alone **sysadm** contains a built–in subset of shell commands that are useful for repairing damaged file systems whereas stand–among **sysadm** has at its disposal a full repertoire of shell commands.
- The software installation menu is unique to stand–alone **sysadm**.

When to Use Stand–among sysadm

By default, use this version, **/usr/sbin/sysadm**, while the DG/UX system is booted and running at init level 1 or higher. If you are not logged in as superuser, some **sysadm** operations will be restricted from use. Restricted operations are shaded in the graphical version and surrounded by brackets in the ASCII version.

When to Use Stand–alone sysadm

Use this version, **/usr/stand/sysadm**,while the DG/UX system is not running. You boot this version to perform operations that you cannot perform while running the installed version of the DG/UX software.

If, for example, the / (root) or **/usr** file systems become damaged and you cannot boot the system, use stand–alone **sysadm** to repair them and then retry booting the system. Refer to Chapter 3 for information on repairing damaged DG/UX file systems. You must take the system down, and then boot stand–alone **sysadm** from the SCM. An example follows:

```
SCM> b  sd(cisc(),0)usr:/stand/sysadm ↵
```

You can boot stand-alone **sysadm** from disk only if the virtual disks containing the **/usr** file system are located on a single physical disk. If **/usr** spans multiple physical disks, you cannot boot from **/usr**.

You may also use stand–alone **sysadm** to copy the physical disk that contains the / file system to another physical disk.

Differences from Previous Releases

- The online storage management (OSM) facility allows you to manipulate your physical disk partitions without having to take your data off line by unmounting file systems. OSM offers a flexible hierarchy that allows you to dynamically restructure configurations. OSM enables on–line backups and on–line data storage failure recovery. VDM, a primary component of OSM, requires that LDM–formatted physical disks be converted to a VDM equivalent.
- This release supports new **rename** and **move** operations and improved **copy** and **expand** virtual disk operations that allow you to manipulate virtual disks while they are online and in use.
- This release supports booting from a multi–pieced (aggregated) virtual disk as long as the pieces reside on the same physical disk.
- In soft formatting, among other changes, surface analysis is no longer recommended and installation of a bootstrap is optional. In addition, bad block remapping is not included with the creation of the virtual disk information table (formerly system areas), and there is support for a new operation to convert between logical and virtual disk formats.
- This release supports use of both the NVRAM board and disk drive as front–end virtual disk cache devices. In addition, it supports front–end device sharing among multiple back ends.
- This release replaces the stand–alone and stand–among versions of **diskman** with stand–alone and stand–among versions of **sysadm**. Refer to a later section in this chapter for more information.
- This release introduces an expanded set of DG/UX shell commands that are built in to stand–alone **sysadm**.

Converting Physical Disks to Virtual–Disk Format

Conversion from logical–disk format to virtual–disk format is necessary only if you are upgrading from a previous release of DG/UX to DG/UX 5.4R3.00. During the upgrade, you have two options for making the transition:

- One–time operation to convert all physical disks from a logical disk format to a virtual disk format.
- Continued use of logical disks in a virtual disk environment (compatibility mode).

Refer to *Installing the DG/UX™ System* for procedures to perform the one–time operation.

The physical disk's metadata format must change to accommodate the virtual disk technology. Metadata describes the layout of a physical disk and how it is partitioned. Also, the new metadata is sufficiently flexible to accommodate future disk formats.

Only the physical disk metadata changes; the user data is not touched. The metadata is converted only after it is safe to do so—without risk of data loss or corruption. Also, there is no risk of format conversion failure because of limited space. The process of converting a physical disk's metadata takes about a second per physical disk.

In most cases, you will want to convert all readable and writable logical disks to virtual disks. Physical disks that contain logical disks **root**, **usr**, and **swap** must be converted; they cannot be used in compatibility mode.

Candidates for compatibility mode are physical disks that may be used in both LDM and VDM environments and read–only devices such as CD–ROM devices. You will be restricted to operations that do not expand, shrink, delete, or re–create the disk's metadata. You may continue to read and write data to physical disks in compatibility mode.

Refer to "Converting a Physical Disk between Logical and Virtual Disk Formats" for information on converting physical disks. Refer to "Registering a Physical Disk" for instructions on registering a physical device in compatibility mode.

Replacing Stand–alone diskman with Stand–alone sysadm

Stand–alone **sysadm** replaces the disk management utility that supported releases prior to DG/UX 5.4R3.00—stand–alone **diskman**. It provides all the **diskman** capabilities, but through the **sysadm** interface. Both stand–alone **diskman** and stand–alone **sysadm** provide a subset of shell commands for administrative operations such as repairing damaged file systems or recovering from system installation problems. Stand–alone **sysadm**, however, offers a significantly greater number of shell commands than its **diskman** predecessor. You access these commands from any stand–alone **sysadm** menu prompt by the shell escape, **!** followed by Enter.

Table 7–1 lists the shell commands in **/sbin** that stand–alone **sysadm** supports:

Table 7–1 Shell commands in /sbin

fsck	init	unmount	sh
halt	mount	reboot	su

Table7–2 lists the shell commands in **/usr/sbin** that stand–alone sysadm supports:

Table 7–2 Shell commands in /usr/sbin

devnm	gridman	probedev	unmount
dg_sysctl	halt	reboot	xdrtoc
exportfs	mkfs	swapon	
fsck	mount	syslogd	

Table7–3 lists the shell commands in **/usr/sbin** that stand–alone **sysadm** supports:

Table 7–3 Shell commands in /usr/bin

admdefault	cut	idi_confirm	rmdir
admdevice	date	idi_doop	sde_target
admfilesystem	dc	idi_echo	sed
admkernel	dd	idi_error	sh
admpackage	df	idi_log	sort
admpdisk	diff	idi_warning	stty
admrelease	dirname	ifconfig	su
admservice	du	kill	sync
admtape	ed	ln	tail
admvdisk	egrep	logger	tar
awk	expr	ls	tee
basename	false	mkdir	touch
cat	find	mt	tput
chgrp	grep	mv	tr
chmod	gunzip	netinit	true
chown	gzip	newaliases	tty
comm	head	ping	uncompress
compress	hostname	pmtd	uniq
cp	id	printf	who
cpio	idc	pwd	xargs

Making a Physical Disk and its Contents Usable

This section applies to normal hard physical disks such as a disk drive or unit in a CLARiiON™ disk–array storage system. For devices such as CD-ROMs, magneto-optical devices, diskette drives, and tape drives see Chapter 15.

You must perform the following steps to get a device recognized in your hardware configuration and to make it accessible.

1. Configure the device into the kernel.

   If you just installed the DG/UX system, the attached devices are already autoconfigured. If you are adding a device to an operational system, you must configure it explicitly before it is recognized.

   Use the **sysadm** operation `Device -> Configure`.

   Go to "Configuring Physical Devices" for details.

2. If the disk is not already soft formatted, do so.

Use the **sysadm** `Device -> Physical -> Soft Format` **menu.**

Go to "Soft Formatting a Physical Disk" for details.

3. If the disk has not already been registered, do so.

   Use the **sysadm** operation `Device -> Physical -> Register`.

   Go to "Registering a Physical Disk" for details.

4. Create a virtual disk on the physical disk to later house data or software.

   Integral to this step is the option to overlay a DG/UX file system on the newly created virtual disk. The file system contains the internal data structures that the operating system requires to keep track of files and directories. If you intend to load DG/UX add–on software or DG/UX data files into the storage area, you must create a file system first. If, however, you intend to load a third–party package, consult its release notice for such requirements.

   There are two general methods for creating a virtual disk:

- You don't care where the virtual disk is located.
  Use this method when you do not care where the space to be used for the virtual disk is located; you do not specify the physical disk(s), the starting blocks, or name another virtual disk as a source for space. You have the option to create a file system on the virtual disk just created. But, most configuration options are limited; in particular, you don't have the capability to stripe.

  If you chose to create a virtual disk without a file system, you have another opportunity to create a file system using the **sysadm** operation `File System -> Local File System -> Create` operation.

  See "Creating File Systems" in chapter 9 for more information.

  Some software applications do not require the use of a file system. Therefore, you may opt not to create one. Refer to your application's documentation for details on its requirements.

- You do care where the virtual disk is located.
  Use this method to control where you get space on a physical disk and to stripe the storage.

  The remaining steps track the requirements for creating a virtual disk, wherein you control its location.

Managing Physical Disks

The tasks you typically perform on physical disks are:

- Configuring and deconfiguring
- Soft formatting
- Registering and deregistering

- Copying
- Listing attributes
- Tracking bad blocks
- Converting a physical disk between logical and virtual disk formats
- Repairing a virtual disk information table

Configuring a Physical Disk

To dynamically configure a physical disk while the system is operating, select the **sysadm** operation `Device -> Configure`. Configuring a physical disk puts an entry for it in **/dev/pdsk** and **/dev/rpdsk**. This operation is useful when physically adding a new device (either temporarily or long term) to a system that is already operating.

To dynamically add a physical disk to your configuration during operation, configure it using the **sysadm** operation `Device -> Configure`. The kernel must already contain the driver for the device being configured. For example, you may configure a SCSI disk in a system in which you have previously configured a SCSI device attached to a integrated SCSI controller—**insc**(). However, if you have not previously configured the driver, for example, **dgsc()**, you must do so before you can configure the device. In this case, you are advised to build a new kernel, including the new driver and attached device(s) with the **sysadm** operation `System -> Kernel -> Build`. See Chapter 4 for information on building a kernel and *Customizing the DG/UX™ System* for information on specifying devices.

By default, at boot time, the system configures all disks that have entries in the system file used to build the kernel. The system file contains entries for any physical disks installed at standard locations when the system file was built.

The configure operation does not update your system file with the added device's name. You must rebuild the kernel to include the added device in the system file to ensure that it gets recognized at boot time.

To avoid having to rebuild a kernel to accommodate newly added devices, you can use an abbreviated device specification in the system file. Instead of specifying the device ID in the device specification, use the asterisk (*) wildcard, which is a pattern–matching character that represents all IDs on that controller. Refer to Chapter 4 for more information on kernel building.

CAUTION: If you have a failover configuration, however, do not use the asterisk convention for device names in your system file.

Deconfiguring a Physical Disk

To deconfigure a physical disk, select the **sysadm** operation `Device -> Deconfigure`.

To make a disk inaccessible, deconfigure it using the **sysadm** operation `Device -> Deconfigure`. The operation requires that you select the desired physical disk's device name, for example, **sd(cisc(0),0,0)**, from a list of configured disks.

If the physical disk is registered, you must deregister it before deconfiguring it. If the physical disk is mounted as a file system (as is typically the case with diskettes), you must unmount it before deregistering and deconfiguring it. Deconfiguring a disk removes its entries from **/dev/pdsk** and **/dev/rpdsk**.

If the system file contains the deconfigured device and you do not want to configure it, you must eventually remove it from the system file and rebuild the kernel. If you used the asterisk (*) notation in the device specification in the system file, you do not have to delete it from the system file and rebuild the kernel.

Soft Formatting a Physical Disk

Formatting a physical disk involves installing or creating these components on the disk:

- Disk label
- Virtual disk information table (VDIT), which was known as system areas in all releases prior to DG/UX 5.4R3.00
- Bootstrap
- Bad block remapping

A physical disk must be deregistered before it can be software formatted.

The following sections discuss these operations.

Installing a Physical Disk's Label

To install a physical disk label, select the **sysadm** operation `Device -> Disk -> Physical -> Soft Format -> Label Disk`.

Disk labels contain the disk geometry (such as tracks per cylinder, bytes per sector, and so on), information that the system requires to write to the disk, read from it, and keep track of damaged disk blocks.

After you select this operation, the system displays the following queries:

`Physical Disk(s): [sd(ncsc(0,7),6,0)]`
: Specify the device on which to install a label.

`Disk Type:`
: Select the type of the disk physical disk to be labeled. The choices are:

1 Model 6442: full–height ESDI, 327 MB
2 Model 6541 or 6542: SMD, 1066 MB
3 Model 6555: full–height ESDI, 648 MB
4 Model 6661: half–height ESDI, 330 MB
5 Generic SCSI

The physical disk is also registered when it is labeled.

An example of a disk label, including the bootstrap, follows:

```
Disk Label:
cylinders_per_drive     0
visible_cylinders_per_drive     0
tracks_per_cylinder     0
sectors_per_track       0
bytes_per_logical_sector        0
bytes_per_unformatted_sector    0
defect_info_start_sector        0
bytes_in_defect_info    0
number_of_relocation_areas      0
sectors_per_relocation_area     0
next_relocation_sector  0
interleave      0
head_skew       0
cylinder_skew   0
head_group_skew 0
spares_per_track        0
bytes_per_data_preamble 0
bytes_per_id_preamble   0
base_head_for_volume    0
flags   0
bytes_in_gap_1  0
bytes_in_gap_2  0
sanity_flag     305441741
version_number  1
Bootstrap: start = 17, size = 500, version = 1
```

To display a disk's label, select the **sysadm** operation `Device -> Disk -> Physical -> List`, which is described in a later section.

Creating a Virtual Disk Information Table (VDIT)

NOTE: This operation replaces the create system areas operation in releases prior to DG/UX 5.4R3.00.

To create a virtual disk information table, select the **sysadm** operation `Device -> Disk -> Physical -> Soft Format -> Create VDIT`.

CAUTION: If a physical disk already contains data, this operation destroys all data on the physical disk.

A virtual disk information table enables the physical disk to receive virtual disks that you create explicitly. It controls the mapping of virtual disks. You must create a virtual disk information table before you create virtual disks. If you intend for an application to directly access a raw physical disk, do not create a virtual disk information table on it.

The physical disk is also registered when its virtual disk information table is created.

An example of a newly created virtual disk information table, including the bad block remap area and bootstrap, appears as follows:

```
Disk name                  State   Reg? Format Total blocks  Free blocks
sd(insc(0),0,0)            avail     y  vdisks      1295922      1230279

   Partition Name                   Role             Address         Size
   .Label,2CA9A8E1                                         0            1
   .Primary_Vdit,2CA9A8DF                                  1           16
   .Bootstrap,2CA893F1                                    17          500
   .Primary_Bad_Block_Table,2CA893F2                     517            5
   .Remap_Area,2CA893F3                                  522          100
   .Secondary_Bad_Block_Table,2CA893F4                   622            5
      <free space>                                     65627      1295279
   .Secondary_Vdit,2CA9A8E0                          1295906           16
```

System partitions begin with a period (.). In this example, roughly the first 620 blocks and the last 16 blocks on the physical disk are reserved for system partitions. The numeric strings appended to the system partitions differentiate them from other system partitions located on other physical disks.

If you are familiar with previous releases of the DG/UX system, you will see that the new virtual disk information table uses for overhead not only the beginning of the physical disk but the final 16 blocks of the physical disk as well. To guard against accidental destruction of a virtual disk information table, the system puts two separate virtual disk information tables on a physical disk.

In some cases, such as to reinstall the operating system, you may want to use this operation to deliberately destroy the data on a physical disk.

Among the system areas that the operation creates are the label and the primary virtual disk information table which arranges the virtual disks that are on the physical disk.

When preparing a diskette, it is not worthwhile to create a virtual disk information table on the diskette. Diskettes are so small that the system areas would use up too much space. Typically, you create a single file system on the entire diskette. For more information on preparing diskettes, see Chapter 15.

Establishing Bad Block Remapping

To establish bad block remapping, select the **sysadm** operation `Device -> Disk -> Physical -> Soft Format -> Mapping`.

NOTE: You do not need to map highly reliable devices, such as RAID mirrors or nonvolatile RAM disks (NVRDs).

The system uses a bad block remap area to store known good blocks to replace blocks that go bad elsewhere on the disk. When you create the virtual disk information table on a disk, the operation installs a bad block remap area and bad block table. The default size is generally sufficient.

A physical disk must be registered before you can establish bad block remapping on it.

After establishing bad block mapping on a physical disk, you may list the blocks that were mapped or unmapped using `Device -> Disk -> Physical -> Bad Blocks`. Go to a later section on tracking bad blocks on a physical disk for more information.

Installing a Bootstrap

To install a bootstrap, select the **sysadm** operation `Device -> Disk -> Physical -> Soft Format -> Install Bootstrap`.

A physical disk must have a current bootstrap if you are to boot from the disk. You are advised, as a matter of course, to install a bootstrap on the physical disk in case, at a later date, one is required. It uses only 500 blocks.

If you are adding a new release of the DG/UX system, you need to reinstall the bootstraps on the disks containing the **root** and **usr** virtual disks. Normally, a new bootstrap is installed automatically during an installation or upgrade. You should install bootstraps on any disk from which you intend to boot, whether booting the kernel (**/dgux**), stand-alone **sysadm** (**/usr/stand/sysadm**), or any other bootable image.

Performing All Soft Formatting Steps

To perform all soft formatting steps, select the **sysadm** operation `Device -> Disk -> Physical -> Soft Format -> All Steps`.

You can select this operation to perform all of the soft formatting steps (install disk label, create virtual disk information table, establish bad block remapping, and install bootstrap).

Registering a Physical Disk

To register a physical disk, select the **sysadm** operation `Device -> Disk -> Physical -> Register`.

You register a physical disk to make its virtual disks accessible on the system. You need to know the physical disk's device name, for example, **sd(dgsc(0),1)**, in order to register it.

You can register a physical disk only if it contains a virtual disk information table, which is used by the system to track virtual disks. You create a virtual disk information table with `Device -> Disk -> Physical -> Soft Format -> Create VDIT`. Diskettes and CD-ROM disks that contain High Sierra or ISO 9660 file systems do not contain a virtual disk information table and do not require registration. For more information on diskette and CD-ROM disk drives, see Chapter 15.

Registering a physical disk makes the physical disk's virtual disks known to the system. You cannot access a virtual disk unless you register the disk drive containing it.

If you register a physical disk that contains logical disks, the physical disk will be registered in compatibility mode. Should you prefer to continue using logical disks in a virtual disk environment, you may do so with restrictions. You can continue to read and write data, but you cannot perform any operations that alter the disk's metadata such as moving, copying, expanding, or shrinking a logical disk or file system. In general, you are advised to convert all physical disks in logical disk format to virtual disk format. Refer to a later section on converting physical disk formats for more information.

If the physical disk does not exist, is not soft formatted, or if either of its node files (in **/dev/pdsk** or **/dev/rpdsk**) is already open, the Register operation will fail. A physical disk's node file may be open if a database management system has opened it for direct access, for example.

Deregistering a Physical Disk

To deregister a physical disk, select the **sysadm** operation `Device -> Disk -> Physical -> Deregister`.

You deregister a physical disk to make its virtual disks (or logical disks, if used in compatibility mode) inaccessible. To deregister a device, you must know its name; for example, **sd(dgsc(0),1)**.

Deregistration fails if any virtual disk on the physical disk is in use; for example:

- A virtual disk's file system is mounted.
- The **swap** virtual disk is located on it.
- An application is using a virtual disk.
- The parent virtual disk for a child virtual disk being used elsewhere is located on the physical disk.

You should deregister a disk before powering it down, or removing it from the configuration. Before removing a CD-ROM or magneto-optical disk from the disk drive, unmount the file systems and deregister the physical disk (if registered). If you deregistered a compact disk, optical disk, or diskette drive, you may remove the disk from the drive without turning off power to the computer or drive.

When using stand–among **sysadm**, you cannot deregister a physical disk that is considered busy or open, such as one that contains any part of the operating system—**root**, **usr**, **swap**, **usr_opt_X11**—or that is being accessed. You must take the system down and use stand–alone **sysadm** instead.

If you deregister a disk to disconnect it from the system, make sure that you turn off its power before removing it. You do not need to turn off power, however, for disks in disk–array subsystems.

Copying a Physical Disk

To copy a physical disk, select the **sysadm** operation `Device -> Disk -> Physical -> Copy`.

The physical disk copy operation copies the contents of a source disk to a destination disk. The contents will be identical, but the location of the content on the destination disk may be different. You can use this operation to copy between disks of different types and sizes.

The **sysadm** Copy operation performs an over-writing copy whereby the contents, if any, of the destination disk are obliterated before the source contents are written to it.

Prerequisites for a Physical Disk Copy

The physical disk must be inaccessible to users for the duration of the copy operation. To take the source physical disk offline you must:

- Unmount all file systems contained on the source disk using the operation File `System -> Local Filesys -> Unmount`. Note the file systems are unmounted so that you can remount them following the copy operation. List the contents of the source disk to help you determine the file systems to be unmounted using the operation `Device -> Disk -> Physical -> List`.
- Register the source disk using the operation `Device -> Disk -> Physical -> Register`.
- Deregister the destination disk using the operation `Device -> Disk -> Physical -> Deregister`.

Why and When to Copy

There are several reasons for copying a physical disk:

- When you replace an old disk with a new one, you can copy the contents of the old one to the new one.
- When installing a number of systems that need the entire operating system loaded from tape, you can save time by loading only the first disk from tape and then copying the other disks from it.

  To use the copy facility for this purpose, you may need to know how to jumper the destination disk for a different address or SCSI ID. This is particularly true if you intend to use the copy facility to load one workstation's SCSI disk from another workstation's SCSI disk. Typically, both SCSI disks come from the factory jumpered as **sd(insc(0),0)**. Before connecting them both to the same system, you need to change the jumpers on one of them to put it at another SCSI ID, for example, **sd(insc(0),1)**.

  If you use the copy facility to make copies of the DG/UX system virtual disks (**root** and **usr**), perform the copy before customizing the DG/UX system. If you perform the copy after setting up, you will duplicate system-specific data (Internet addresses, system names, and so on) from the source disk.
- When you need to remove a disk from your system, use the Copy operation to copy the contents of the disk onto another disk having the required available space.

There may also be other scenarios where you find the physical disk copy operation useful.

Making the Copied Virtual Disk Usable

Following a successful copy, the source and destination disks will have identical names and identical content. Do not register both the source and destination virtual disks. Register only the virtual disk that you intend to use with the **sysadm** `Register` operation. Bad blocks will be not be remapped to the destination. To recover bad blocks for the destination device, use the **sysadm** `Unmap` operation, which is covered in a later section. Finally, mount the destination's file systems using the operation `File System -> Local Filesys -> Mount`, which is covered in Chapter 9.

NOTE: Unlike releases prior to DG/UX 5.4R3.00, the current release allows you to register both the source and destination virtual disks, but doing so produces a virtual disk name conflict. When you register the first virtual disk, both its long and short virtual disk names will appear in the files **/dev/dsk** and **/dev/rdsk**. For example:

```
/dev/dsk/root
/dev/dsk/vdm(root,2CBACF45,0C3B220A,0)
```

When you register the second virtual disk, the short names will conflict and only the fully qualified long name will appear in **/dev/dsk** and **/dev/rdsk**. Therefore, you may access the second copy using its long name only.

Listing Physical Disks

To list statistics about a physical disk, select the **sysadm** operation `Device -> Disk -> Physical -> List`.

After you select this operation, the system displays the following queries:

`Physical Disk(s) [all]:`
: Select the physical disk(s) whose statistics you want to list. You may select all (registered and unregistered devices), only registered devices, or explicit devices by name.

`Listing Style: [normal] :`
: Specify the type of list you want. The List operation provides three list options: normal, partition, and verbose.

List Label? [no]
: Choose this option to list the contents of the physical disk label of each physical disk that is listed. The label describes the geometry of the disk that the operating system uses to manage the space on the physical disk.

A Normal list displays the device name(s), whether or not it is available (not in use by another host such as with dual–ported hosts), and whether or not the disk is

registered. Also, it specifies whether it is formatted for logical disks or virtual disks, its total size in 512-byte blocks, and if registered, the number of free blocks remaining on the disk. An example of a normal listing follows:

```
Disk name                    State   Reg? Format Total blocks  Free blocks
sd(insc(0),0,0)              avail      y vdisks      1295922      1230279
```

An example of a normal listing and its label follows:

```
Disk name                    State   Reg? Format Total blocks  Free blocks
sd(dgsc(0,6),0,6)            avail      y vdisks      2311043      2298977
Disk Label:
cylinders_per_drive     1224
visible_cylinders_per_drive     1219
tracks_per_cylinder     15
sectors_per_track       47
bytes_per_logical_sector        512
bytes_per_unformatted_sector    0
defect_info_start_sector        0
bytes_in_defect_info    0
number_of_relocation_areas      0
sectors_per_relocation_area     0
next_relocation_sector  0
interleave      1
head_skew       3
cylinder_skew   13
head_group_skew 0
spares_per_track        1
bytes_per_data_preamble 0
bytes_per_id_preamble   0
base_head_for_volume    0
flags   0
bytes_in_gap_1  0
bytes_in_gap_2  0
sanity_flag     305441741
version_number  1

Bootstrap: start = 8, size = 500, version = 1
```

A Partitions list provides an enumeration of the user–created virtual disk partitions and system partitions on the physical disk. The listing includes the partition's name, its starting address, and size in 512-byte blocks. In cases where there are multiple partitions forming an aggregation, each partition is identified by its role—a piece of an aggregation and its number in a series. A partition listing represents the information about the virtual disk information table, bad block map area, and bootstrap as `<Various System Partitions>`.

```
Disk name                    State   Reg? Format Total blocks  Free blocks
sd(insc(0),0,0)              avail     y  vdisks      1295922      1230279

   Partition Name                   Role              Address         Size
   <Various System Partitions>                             0          627
   payroll                                               627         5000
   Part of zenith1                  Piece  1 of  2      5627        30000
```

```
   Part of zenith1                  Piece  2 of  2      35627        30000
   <free space>                                         65627      1230279
   <Various System Partitions>                        1295906           16
```

A Verbose list produces a detailed listing of the system partitions and the user–created partitions.

```
Disk name                    State   Reg? Format Total blocks  Free blocks
sd(insc(0),0,0)              avail      y vdisks      1295922      1230279

   Partition Name                   Role              Address         Size
   .Label,2CA9A8E1                                          0            1
   .Primary_Vdit,2CA9A8DF                                   1           16
   .Bootstrap,2CA893F1                                     17          500
   .Primary_Bad_Block_Table,2CA893F2                      517            5
   .Remap_Area,2CA893F3                                   522          100
   .Secondary_Bad_Block_Table,2CA893F4                    622            5
   payroll                                                627         5000
   Part of zenith1                  piece  1 of  2       5627        30000
   Part of zenith1                  Piece  2 of  2      35627        30000
   <free space>                                         65627      1230279
   .Secondary_Vdit,2CA9A8E0                           1295906           16
```

Tracking Bad Blocks on a Physical Disk

To track bad blocks on a physical disk, select the **sysadm** operation `Device -> Disk -> Physical -> Bad Blocks`.

Disk units occasionally develop flaws in the disk surface. Most disk hardware keeps track of these bad parts without depending on the operating system to do it for them; nevertheless, the DG/UX system offers its own bad block tracking mechanism in case the disk hardware fails to detect or remap a flawed disk block.

When the DG/UX system detects a flaw on a disk, it flags the block (a 512-byte portion of disk space) as bad. If a write operation was performed, DG/UX finds a good block to replace the bad block. The operating system takes care of redirecting reads and writes intended for the bad block so that they go to the replacement block instead. A part of the disk called the bad block remap area contains good blocks reserved specifically for this purpose: to replace blocks that go bad elsewhere on the disk.

When the system creates the bad block remap area, it offers a default remap area size. You may specify another size if you like, but the default size is generally sufficient.

Mapping Bad Blocks

To map bad blocks on a physical disk, select the **sysadm** operation `Device -> Disk -> Physical -> Bad Blocks -> Map`.

Use this operation to add bad blocks to a physical disk's bad block table. Bad blocks are parts of the physical disk that may be unreliable for storage and retrieval of

information. The operating system keeps track of bad blocks by listing them in its bad block table. There may be other bad blocks besides the ones listed in the bad block table. If you suspect that a particular block is unreliable or if diagnostics have shown a block to be unreliable, you can add that block to the bad block table using the mapping operation.

To add a bad block to the bad block table, you need to know the device name of the physical disk and the physical address of the bad disk block. The disk should be registered before you invoke the operation. You do not need to know the address of the replacement block.

If you map a block that contains data, the system will not attempt to copy the data from the bad block to the replacement block. The contents of the replacement block are unknown until a write operation writes data to the block.

If you tell DG/UX to force remap a block, it does so without verification of the original block's status, and without confirmation. A forced block ends up as mapped after it is written.

If the bad block remap area does not have any more good blocks (which is highly unlikely), the operation returns an error message telling you that it cannot add the bad blocks to the table. When this happens, back up the disk contents and re–create the system areas on the disk, being sure to specify a larger bad block remap area size. Then reload the disk from backups.

After you select this operation, the system displays the following queries:

`Physical Disk:`
: Specify the name of the physical disk for which you wish to map bad blocks.

`Block Numbers:`
: Enter the number of one or more blocks to map. Enter a block number, a list of block numbers (separated by commas), or a range of block numbers (two numbers separated by a hyphen).

  After you confirm your desire to perform the operation, the remapping is performed. The system displays messages that confirm the remapping.

Unmapping Bad Blocks

To unmap (recover) bad blocks on a physical disk, select the **sysadm** operation `Device -> Disk -> Physical -> Bad Blocks -> Unmap`.

The operation recovers the remapped blocks that no longer need to be remapped. To remap a block means to designate another block to use in place of the bad or unreliable block. For example, block 100000 might be remapped to another block. After getting your disk serviced, however, you may want to remove block 100000 and all other formerly bad blocks from the bad block table.

After you select this operation, the system displays the following queries:

`Physical Disk:`
: Specify the name of the physical disk for which you wish to unmap bad blocks.

`Block Numbers:`

Enter the number of one or more blocks to unmap. Enter a block number, a list of block numbers (separated by commas), or a range of block numbers (two numbers separated by a hyphen).

After you confirm your desire to perform the operation, the recovery is performed. The system displays messages that confirm the recovery of blocks.

Listing Bad Blocks

To list bad blocks on a physical disk, select the **sysadm** operation `Device -> Disk -> Physical -> Bad Blocks -> List`.

After you have formatted a disk, use this operation to display bad blocks (if any). If you supply the disk drive name, the operation lists the addresses (in decimal form) where bad blocks are located.

After you select this operation, the system displays the following query:

`Physical Disk:`

Specify the name of the physical disk whose bad block table you wish to view.

An example of a bad block table follows:

```
Physical Disk(s): sd(ncsc(0,7),0,0)
       Block  Index  Status
          20      0  force
          21      1  force
          22      2  force
          23      3  force
          24      4  force
          25      5  force
          26      6  force
```

The first column lists the bad blocks that you specified for remapping. The second column identifies the block number in the remap area to which the bad block is remapped. The third column specifies the status of the bad block remap area.

`mapped`

The system writes the bad block to a new block. All I/O is being redirected to this new block.

`unmapped`

The kernel detects an error upon attempting to read this block. Although the original block has a remap block associated with it, the contents of the remap block are undetermined. Until a write is performed to the new block, reads directed to the original block will fail.

`force`

A user performs a a map–block operation on this block. Although the original block has a remap block associated with it, the contents of the remap block are undetermined. Until a write is performed to the new block, reads directed to the original block will fail.

`pseudo`

An entry with this status corresponds to a block that does not necessarily need remapping but whose contents are undetermined. Reads to pseudo bad blocks will fail. The first write to a pseudo bad block sets its contents and deletes the corresponding remap table entry.

`bad`

An entry with this status represents a block in the bad–block map area itself that is unusable.

Converting a Physical Disk Between Logical and Virtual Disk Formats

To convert a physical disk between logical and virtual disk formats, select the **sysadm** operation `Device -> Disk -> Physical -> Convert`.

We strongly encourage you to convert physical disks from logical disk format to virtual disk when upgrading the DG/UX system.

Use this operation to convert the entire physical disk from logical disk format to virtual disk format. After such a conversion from logical to virtual disk format, a physical disk can revert back to logical disk format. Physical disks that were originally formatted for virtual disks, however, cannot be converted to a logical disk format.

Only the physical disk metadata changes; the user data is not touched. The metadata is converted only after it is safe to do so—without risk of data loss or corruption.

The process of converting a physical disk's metadata from logical disk to virtual disk format takes about a second per disk and does not endanger the data. Also, there is no risk of format conversion failure because of limited space.

Why Convert?

The physical disk's metadata format must change to accommodate the virtual disk technology. Metadata describes the layout of a physical disk and how it is partitioned. Also, the new metadata is sufficiently flexible to accommodate future disk formats.

Conversion Prerequisites

Note the following requirements when converting physical disks from logical disk format to virtual disk format:

- All physical disks to be converted must be configured, powered on, writable, and deregistered.
- All physical disks that contain the DG/UX operating system — **root**, **usr**, and **swap**—must be converted to virtual disk format during the upgrade to DG/UX 5.4R3.00.

- Read–only devices such as CD–ROM drives or WORM (write once, read many) drives cannot be converted to virtual disk format.
- All pieces of a multi–pieced logical disk should be converted at the same time to maintain the relationships of components that span physical disks. Failure to do so renders the incomplete virtual disk unusable. To recover, you may assemble errant components by hand. Refer to the next section for a discussion of rebuilding virtual disks by hand.
- All mirror images should be converted at once to maintain the relationships of components that span physical disks. Failure to do so may render the incomplete mirror virtual disk unusable. Success or failure depends on the maximum number of lost images to be tolerated (see the instructions for mirroring a virtual disk). To recover, you may assemble errant images by hand. Refer to the next section for a discussion of rebuilding virtual disks by hand.
- Caches must be dismantled (deleted) before their host physical disks can be converted.

Some virtual disk hierarchical constructs may not successfully revert to logical disk format because of hierarchical complexity. For example, conversion of an aggregated virtual disk composed of other aggregations to the logical disk equivalent does not work. If a backward conversion fails, your virtual disks remain intact. You may continue to use simple partitions, aggregations, and mirrors; but you cannot use higher order configurations such as aggregated aggregations or aggregated mirrors.

After you select this operation, the system displays the following queries:

`Conversion target format: [Virtual disk format]`
: Specify the format to which you want to convert the physical disk.

`Forcefulness: [Careful]`
: Select the manner in which incomplete conversions will be handled. If you have a virtual (or logical) disk that spans multiple physical disks, you should convert all affected physical disks in one operation. If you convert a subset of the affected physical disks, the conversion will produce disconnected virtual disk components that will require you to rebuild the virtual (or logical) disk by hand. Refer to the next section for a discussion of rebuilding virtual disks by hand.

  `Forceful` Despite any problems, if any, encountered by the conversion program, the conversion will proceed, and will produce incomplete and unusable virtual (or logical) disk components.

  `Careful` If the conversion program encounters a problem, the conversion will be aborted.

  `No-write` The conversion program performs a trial conversion. It reports any problems, if encountered, but will not convert the physical disk(s) regardless of success or failure.

  You are advised to select either the careful or no–write option before you select forceful.

`Physical disks to convert`

Specify the DG/UX device names of the physical disks to convert.

Building Virtual Disks by Hand from a Failed Conversion

If a conversion to virtual disk format fails, the "orphan" components will remain intact and unusable. You may build the desired virtual disk "by hand" using the unconverted components only by correctly identifying those errant, unconverted components. You may do so by listing the contents of physical disks and recognizing unnamed children and their sizes. Such an endeavor will depend largely on your knowledge of locations and sizes of logical disks on physical disks in your configuration. This can be a difficult task. Still, it may be possible to correctly link several unnamed, same–sized mirror images to a new virtual disk mirror or to create an aggregation from several unnamed partitions.

Repairing a Damaged Virtual Disk Information Table

To reconstruct a damaged virtual disk information table (VDIT) on a physical disk, using the duplicate that was created when the virtual disk information table was created, select the **sysadm** operation `Device -> Disk -> Physical -> Repair VDIT`. Damage can result, for example, when you attempt to boot a pre–DG/UX 5.4R3.00 kernel on a DG/UX 5.4R3.00 system.

Use this operation if you see on the console a message such as:

```
Error:  The VDIT for disk 'sd(insc@7(FFF8A000,7),1,0)' has
        degraded and is using only one VDIT copy.
```

The physical disk must be deregistered and reregistered before the VDIT repair can take effect. The repair operation will attempt to deregister and reregister the physical device; but, will fail if it can't deregister the device. If this happens, you may continue to use the disk in degraded mode until you are able to take the system down to deregister and reregister the physical disk. A degraded mode of operation allows read and write operations to continue, but prohibits operations that alter the disk's metadata such as enlarging, shrinking, moving, and mirroring. While operating in a degraded mode is allowed, you are discouraged from doing so for an extended period.

After you select this operation, the system displays the following query:

`Physical Disk:`

Specify the name of the physical disk that contains the virtual disk information table.

Managing Virtual Disks

A virtual disk is an amount of space that you reserve on a physical disk onto which data is stored. A virtual disk can span multiple physical disks. A virtual disk is further characterized by a type: partition, aggregation, mirror, cache, or a combination of virtual disk types. Refer to the beginning of the chapter for a review of these terms and definitions.

Creating a Virtual Disk

Before you actually create a virtual disk, make sure you have performed the prerequisite operations. Refer to "Making a Physical Disk and its Contents Usable" for a review.

If you want to create a virtual disk quickly and you do not care what physical disk it is on, where on the physical disk it is, and you do not care to stripe the data, you may choose to use the **sysadm** `File System -> Local FileSys -> Create` operation. See "Creating File Systems" in chapter 9 for more information.

If you want control over where the virtual disk is located, want striping, and other configuration options, use the **sysadm** `Disk -> Virtual -> Create` operation.

After you select the create operation, the system prompts for various attributes.

Virtual Disk Name

A virtual disk name may be no more than 31 characters in length. The following table lists the characters to be excluded from a virtual disk name.

Table 7–4 Illegal Characters in Virtual Disk Names

ISO–8859 Character	Description
\000 through \037	ASCII control characters
	space
"	double quote
'	single quote
(	left parenthesis
)	right parenthesis
,	comma
/	slash
:	colon
@	at sign
\177	ASCII DEL
\200 through \237	undefined characters

You may use all other characters in the ISO–8859 character set in virtual disk names. You may name your virtual disks according to a function or classification

that fits your application, such as a virtual disk named **tax_86** that contains tax records for 1986.

After you select `Disk -> Virtual -> Create`, the system displays this prompt:

`New Virtual Disk Name`

Enter the desired name for the virtual disk being created.

Should I Stripe?

After you supply the name of the virtual disk, **sysadm** then asks whether or not you want to stripe it. Striping is an optional attribute of an aggregated virtual disk (one typically formed by combining two or more identically sized partitions).

Depending on the nature of your applications, you may find that data striping improves disk I/O performance. You can implement data striping through the hardware (if you have a disk-array subsystem) or through the software. This manual covers only software data striping. For information on hardware data striping, see your disk-array documentation.

To implement software data striping, you need to create the virtual disk with this purpose in mind. Once you have created the striped virtual disk and its file system, striping is transparent to your applications. You use and manage the striped file system just like any other file system. The only difference is that you cannot change the size of a striped virtual disk; you cannot expand or shrink it.

Striping is implemented by placing consecutive file elements in the file system so that they alternate from one partition of the virtual disk to the next. For example, in the case of a striped aggregation, each partition must be the same size and each should reside on a different physical disk. The system will place the first data element in the first partition, the second data element in the second partition, and the third data element in the third partition, and so on. The data elements alternate this way so that consecutive data elements are stored on alternating disks. The performance advantage results not only because you have distributed the I/O load across three disks, but also because you are using the hardware's read-ahead implementation to get the next element on that disk, even before you have explicitly requested it.

Before continuing, you need to consider whether or not striping can benefit your application's performance. First, striping is only possible if you have more than one physical disk. Second, striping only helps applications that perform a lot of random reads and writes or a lot of sequential reads. Data striping does not help applications that perform a lot of sequential writes.

If your application does not appear suited to striping, do not attempt to implement striping: striping can have a negative impact on performance for inappropriate applications.

The virtual disk that you intend to create for striping must conform to these guidelines:

- The aggregated virtual disk must consist of multiple components (e.g., partitions or other aggregations).
- Each component should reside on a different physical disk.

- Each component forming the aggregation must be the same size, and the size of each component must be evenly divisible by the stripe size. For example, possible stripe sizes for a 400–block partition are 2, 4, 5, 8, 16, 25, 50, 100, and 200. If you intend to put DG/UX file systems on the virtual disks, for maximum performance, set the stripe size to be a multiple of the data element size. The default data element size is 16 blocks. A virtual disk component size must be a multiple of the stripe size.

  If you are not creating a DG/UX file system on each component, select a stripe size that would be most beneficial for the application. A suitable stripe size is generally a multiple of the data element size. Good choices for stripe size are 16, 32, and 64, for example. See the **mkfs**(1M) manual page for more help on stripe size.

- You cannot stripe an existing file system or virtual disk. You may stripe a virtual disk only when you create it.

- You may not stripe the root or **/usr** file systems, or the **swap** virtual disk.

The Striping Procedure

After you supply the name of the virtual disk, **sysadm** then asks whether or not you want to stripe it. If you choose to stripe it, the operation prompts for the stripe size in blocks.

The next prompt follows:

`Striped? [no]`
: Choose whether or not to stripe the virtual disk.

`Stripe Size (in blocks): [16]`
: The stripe size is the number of blocks that are contained on a stripe. The stripe size must be an integral divisor of the number of blocks in each stripe. Refer to the previous section that discusses the striping concepts.

After you have created the striped virtual disk and file system, you may use the file system just like any other file system.

To list information about a striped virtual disk, or to see if a virtual disk is striped, use the **sysadm** operation `Device -> Disk -> Virtual -> List`.

Creating a File System

If you intend to load DG/UX add–on software or DG/UX data files into the storage area, you must create a file system first. If, however, you intend to load a third–party package, consult its release notice for such requirements. A file system contains the internal data structures that the operating system requires to keep track of files and directories. Creation of a file system does not create an entry in the **/etc/fstab** file, nor does it mount and export the file system. After you create a virtual disk, you must add a file system using the `File System -> Local Filesys` menu, which is discussed later in this chapter.

You do not want to create a file system on a virtual disk to be used as a secondary mirror image. However, you will want to create a file system on the virtual disk to be used as the primary image. See a later section on software mirroring for more details.

The **mkfs** utility is what the system uses to create a file system. See the **mkfs**(1M) manual page for more information.

CAUTION: Creating a file system destroys any data that already may be on the virtual disk.

The next prompt follows:

```
Create File System?  [yes]
```

Decide whether or not to put a file system on the virtual disk being created.

Deciding on Size

You specify storage in 512–byte blocks. When you select the size for a virtual disk, it is important to keep in mind that not all of the space in the virtual disk will be available for storing files. File system overhead, in the form of internal data structures, will consume some of the virtual disk space, thus making it unavailable to you. A DG/UX file system by default also has a reserved free space buffer (10% by default) for which you must plan. This space is reserved for the superuser. Assuming that you create a DG/UX file system with the default characteristics, you should create the virtual disk 17% larger than the amount of data that you intend to store in the file system.

For example, if you need a DG/UX file system to hold 100 Mbytes of data, create a virtual disk 117 Mbytes in size. When the file system contains no files, it is around 4% full. After loading it with 100 Mbytes of data, it is 90% full. When a file system is 90% full, only the superuser can extend or create files or directories in the file system.

If you decide at any time that you have created a virtual disk too small, you can aggregate more free space. Keep in mind that if you intend to boot from an aggregated virtual disk, its components cannot span multiple physical disks. Furthermore, to expand or shrink a virtual disk that does not contain a file system, use **sysadm**'s Expand and Shrink operations, which are located in the `Device -> Disk -> Virtual` menu. To expand or shrink a virtual disk that does contain a file system, use **sysadm**'s Expand and Shrink operations, which are located in the `File System -> Local Filesys` menu.

Methods for Creating a Virtual Disk

Sysadm allows you to create a virtual disk using three methods:

- By size alone
- Creating one or more partitions dynamically
- Using one or more existing virtual disks

When creating a virtual disk, you either do care or you don't care where it comes from. If you don't care, you can simply specify the desired size in blocks and **sysadm** locates free space on one or more physical disks. Keep in mind that you cannot stripe a virtual disk that is allocated by size alone. If you intend to stripe, you do care about the locations for virtual disks.

If you prefer to know exactly the location of a virtual disk, you either can collect space directly from the physical disk, partition (carve out) a desired amount from an existing virtual disk, or combine an existing virtual disk with another amount of space. **Sysadm** allows you to create virtual disks using a combination of these methods until you get the amount of space you need.

When starting out with a clean physical disk, your task may be as simple as creating new partitions of a particular size directly from specific physical disks. However, over time, you may choose to organize your resources to achieve a particular performance or high availability standard. These virtual disk creation methods give you the flexibility to arrange and rearrange resources, taking advantage of striping, mirroring, and caching. For example, if your performance goal is to achieve fast I/O, you may initially set up striped disks that span multiple physical disks. Later, as a safeguard against disk crashes, you may decide to mirror those striped disks. In this case, you would want to create mirror images on physical disks different from those housing the striped partitions.

Each virtual disk creation option is described as follows:

Size Alone

Sysadm prompts for the method to use for creating a virtual disk.

```
Select space by:  [Size alone]
```

Select "`size alone`." Or type a question mark (?) for a list of choices. You may type the number that corresponds to this choice, 1.

```
Size in blocks: (1-59795)
```

Enter the desired size. The default range provided varies according to the size of the physical disk.

Depending on the size requested, **sysadm** locates it on one or multiple physical disks. The system then creates the virtual disk, assigning the given name, and making it a volume. Making a virtual disk a volume creates an entry in **/dev/dsk** and **/dev/rdsk**. With a virtual disk's entry in these directories, you are allowed to mount it as a file system or open it to use as a database.

New Partition(s)

Sysadm prompts for the method to use for creating a virtual disk.

```
Select space by:  [Disk to partition and partition size]
```

Select "`Disk to partition and partition size`." Or type a question mark (?) for a list of choices. You may type the number that corresponds to this choice, 2.

`Disk to Partition From:`

Sysadm asks for the disk from which to create the new partition. You can choose space directly from a physical disk's free space or from existing virtual disks. If a virtual disk already exists and is not already open (already used), you may partition space from it. You may partition the entire virtual disk or select a certain number of contiguous blocks from it. Typing ? produces a list of possible sources of space.

Enter the name of the physical disk or the existing virtual disk from which to create a new partition.

`Length of Piece in Blocks: (1-2000)`

Enter the desired size. Pressing Enter will not select the maximum size.

`Starting Block (optional):`

If you requested storage directly from a physical disk, you may select an absolute physical disk starting location. Or, if you requested storage from an existing virtual disk, this number does not correspond to an absolute physical disk location, but rather to a location on the virtual disk itself. Type question mark (?) for the allowable address range. If you do not specify a starting address, the system chooses an appropriate value.

For a disk–array subsystem, to obtain reasonable performance, an I/O operation should start on the hardware stripe that was set when the disk was bound. Align the virtual disk with the hardware stripe by rounding up the virtual disk's starting block number. For example, for a disk–array subsystem whose hardware stripe is bound at 128 sectors, select a virtual disk's starting block that is an even multiple of 128. A legal virtual disk starting block address might be 128, 246, or 384.

`Do you want to specify more pieces for this virtual disk? [no]`

If you do not need any more space, press Enter. Otherwise, answer "yes," and the first prompt in the series repeats. You are not restricted to creating another new partition. You may create storage using any of the allowable methods.

`Virtual disk [`*`virtual-disk-name`*`] is not partitioned. Partition it? [yes]`

This prompt is displayed only if you named an existing virtual disk from which to partition (see second prompt in this section). Since this virtual disk has not been partitioned previously, **sysadm** confirms your desire to do so now. To confirm the request, answer "yes"; otherwise, answer "no" to change your mind.

When you finish, the system then creates the virtual disk, assigning the given name, and making it a volume. Making a virtual disk a volume creates an entry in **/dev/dsk** and **/dev/rdsk**. With a virtual disk's entry in these directories, you are allowed to mount it as a file system or open it to use as a database.

Existing Virtual Disk

Sysadm prompts for the method to use for creating a virtual disk.

```
Select space by:  [Name of existing virtual disk]
```

Select "Name of existing virtual disk." Or type a question mark (?) for a list of choices. You may type the number that corresponds to this choice, 3.

```
Child virtual disk:
```

Sysadm asks for the name of the existing virtual disk to use as a component in an aggregation. Type a question mark (?) for a list of available virtual disks from which to choose. You may type the number that corresponds to the desired virtual disk. Note that you use the child virtual disk entirely.

Enter the name of the existing virtual disk whose space you want to collect.

```
Do you want to specify more pieces for this virtual disk? [no]
```

If you do not need any more space, press Enter to accept the "no" default. Answering "yes" allows you to combine one or more virtual disks with the virtual disk just captured for use. The first prompt in the series repeats. You are not restricted to creating another new virtual disk. You may create storage using any of the allowable methods.

CAUTION: Do not create an aggregation by combining an existing aggregation with any of its child components. A child can be claimed by only one aggregation.

After you finish claiming space for storage (you answer "no" to the "Do you want to specify more pieces for this virtual disk?" prompt), the system identifies the "child" component(s), size, and locations of the new "parent" virtual disk, assigns the given name, and makes it a volume. Making a virtual disk a volume creates an entry in **/dev/dsk** and **/dev/rdsk.** With a virtual disk's entry in these directories, you are allowed to mount it as a file system or open it to use as a database.

Creating Virtual Disks: A Typical Scenario

The following example shows a typical **sysadm** session for creating virtual disks. Consider a configuration where you have three physical disks, one slow and two fast. You want storage on a physical disk for some data that has no fast I/O requirements, and you have a database with critical data for which you need I/O processing speed.

Creating a New Partition

To create the virtual disk that has no fast I/O requirements, perform the **sysadm** `Device -> Disk -> Virtual -> Create` operation.

```
New Virtual Disk Name: dailies↵
Striped? [no]↵
Create File System? [yes]↵
Select Space by: [Size alone] Disk to partition and partition size↵
Disk to Partition From: sd(insc(0),0,0)↵
Length of Piece in Blocks: (1-1295279) 5000↵
Starting Block (optional):↵
Do you want to specify more pieces for this virtual disk? [no] ↵
Virtual disk "dailies" created.
Virtual disk "dailies" made a volume.
File system created on virtual disk /dev/dsk/dailies
```

Creating an Aggregation from Dynamically Created Partitions

Next, to create a virtual disk for critical data that requires fast I/O, you decide to stripe a virtual disk. Striping requires that you use identically sized virtual disks that may span multiple physical disks for improved data distribution. You decide to form a striped aggregation, using two dynamically created partitions (they do not exist yet) on two fast disks, **sd(insc(0),1,0)** and **sd(insc(0),5,0)**. Creating an aggregation in this way produces two unnamed child partitions.

From the **sysadm** `Device -> Disk -> Virtual` menu, create an aggregated virtual disk comprised of two new partitions as follows:

```
New Virtual Disk Name: zenith↵
Striped? [no] yes↵
Stripe Size (in blocks): [16]↵
Create file system? [yes]↵
Select Space by: [Disk to partition and partition size]↵
Disk to Partition From: sd(insc(0),1,0)↵
Length of Piece in Blocks: (1-1290279) 30000↵
Starting Block (optional):↵
Do you want to specify more pieces for this virtual disk? [no] yes ↵
Select Space by: [Disk to partition and partition size]↵
Disk to Partition From: sd(insc(0),5,0)↵
Length of Piece in Blocks: (1-1290279) 30000↵
Starting Block (optional):↵
Do you want to specify more pieces for this virtual disk? [no]↵
30000-block unnamed child partition created at 5627 on "sd(insc(0),1,0)"
30000-block unnamed child partition created at 35627 on "sd(insc(0),5,0)"
Virtual disk "zenith" created.
Virtual disk "zenith" made a volume.
File system created on /dev/dsk/zenith
```

You created a 5000–block partition named **dailies** on slow physical disk **sd(insc(0),0,0)**, and a striped aggregation named **zenith** that contains two

unnamed partitions on fast disks **sd(insc(0)1,0)** and **sd(insc(0),5,0)**. Each partition is 30,000 blocks.

Note that the second set of partitions created was unnamed. In general, having unnamed partitions does not pose a problem as long as you access the parent aggregated virtual disk. However, to access directly a child virtual disk, you must use its long name, represented as *xy*, where *x* is a generation number, and *y* is the system identifier. You may list unnamed child partitions using the list operation, which is covered in a later section. A virtual disk must be named before it can be manipulated. You can assign a name to an unnamed virtual disk in the **sysadm** operation `Device -> Disk -> Virtual -> Rename`. See a later section on renaming virtual disks for details.

Creating a Striped Aggregation from Two Existing Partitions

Instead of aggregating two new partitions, you can create the partitions separately, and then collect them in a striped aggregation. An advantage of this method is that the child partitions are named.

First, you create two separate partitions and then collect them into a striped aggregation.

From the **sysadm** `Device -> Disk -> Virtual` menu, create storage for the first partition as follows:

```
New Virtual Disk Name: part1↵
Striped? [no] ↵
Stripe size (in blocks): [16]↵
Create file system? [no]↵
Select Space by: [Disk to partition and partition size]↵
Disk to Partition From: sd(insc(0),1,0)↵
Length of Piece in Blocks: (1-1290279) 30000↵
Starting Block (optional):↵
Do you want to specify more pieces for this virtual disk? [no]↵
Virtual disk "part1" created.
Virtual disk "part1" made a volume.
File system created on /dev/dsk/part1
```

Repeat the procedure, this time creating a virtual disk named **part2** on the other fast disk, **sd(insc(0),5,0)**.

Next, create a striped aggregation named **zenith**, and select existing virtual disks, **part1** and **part2**, as its children.

```
New virtual disk name:  zenith↵
Striped? [no] yes ↵
Stripe Size (in blocks): [16] ↵
```

NOTE: The next prompt asks if you want to create a file system on the new aggregation. When aggregating existing partitions that already contain

file systems, this operation will overlay the partitions with a new file system, destroying any existing data. If, however, only a file system with no data already exists on the aggregated partitions, no harm is done.

```
Create File System? [yes]↵
Select Space by: [Disk to partition and partition size] Name of an existing
virtual disk↵
Child Virtual Disk: part1↵
Do you want to specify more pieces for this virtual disk? [no] yes↵
Select Space by: [Disk to partition and partition size] Name of an existing
virtual disk↵
Child Virtual Disk: part2↵
Do you want to specify more pieces for this virtual disk? [no]↵
Virtual disk "zenith" created.
Virtual disk "zenith" made a volume.
File system created on /dev/dsk/zenith
```

You created a 60,000–block, striped aggregation virtual disk named **zenith** that contains two existing partitions; **part1** on **sd(insc(0)1,0)** and **part2** on **sd(insc(0),5,0)**.

Removing a Virtual Disk

To remove a virtual disk, select the **sysadm** operation `Device -> Disk -> Virtual -> Remove`. A virtual disk to be removed must be the highest–level (or "parent") virtual disk. Removal of the top–level virtual disk deletes its children virtual disks automatically. If you attempt to remove a "child," such as **part1** in the preceding illustration, the operation will fail. Also, the virtual disk's file system must be unmounted before the virtual disk can be removed.

After you select the **sysadm** operation `Device -> Disk -> Virtual -> Remove` operation, you are prompted:

```
Virtual Disk(s):
```

You are warned of the consequence:

```
Caution:  this operation will destroy virtual disk(s)
<virtual-disk-name>.  Continue? [yes]
```

You cannot remove a virtual disk that is currently in use as a mirror image. There are two ways to remove a mirror image: 1) unlinking an image from the mirror before removing the virtual disk, and 2) dismantling all images from a mirror before removing the virtual disk. Refer to the section on managing mirrors for more information.

Renaming a Virtual Disk

To rename a virtual disk (including system partitions and virtual disks represented with a long name form), select the **sysadm** operation `Device -> Disk -> Virtual -> Rename`. The virtual disk's file system may be mounted and in use when being renamed. The system displays the following prompts:

`Virtual Disk to Rename:`
Supply the name of the virtual disk to be renamed.

`New Virtual Disk Name:`
Supply the name you want to change it to.

After you respond to the system's confirmation, the action is performed automatically.

NOTE: The new virtual disk name is not transmitted automatically to the **/etc/fstab** file. You must specify a different mount point in the **fstab** file using the **sysadm** operation `File System -> Local Filesys -> Modify`.

Expanding a Virtual Disk

To increase the size of a virtual disk that does not contain a DG/UX file system, select the **sysadm** operation `Device -> Disk -> Virtual -> Expand`. To expand a virtual disk containing a DG/UX file system, use the `File system -> Local Filesys -> Expand` operation instead.

Use this operation to expand only partitions and aggregations of partitions. If expanding a partition and sufficient adjacent space exists for the expansion, the partition will subsume the adjacent space. Otherwise, a second partition is created elsewhere to form an aggregation. If expanding an aggregation, you will be creating another partition to include in the aggregation. You cannot expand virtual disks that are striped. To expand a mirror image, refer to the section on mirrored virtual disks.

The system displays the following queries:

`Virtual Disk:`
Select the virtual disk that you want to expand. Use ? for a list of virtual disks to choose from.

`Select Space by: [Size alone]`
You expand a virtual disk by the same methods that you create a virtual disk. Type ? for choices. Refer to the section on methods for creating storage for a discussion.

After you confirm your desire to expand, the operation is performed automatically.

Striped and Software-Mirrored Virtual Disks and Cached Virtual Disks
To change the size of a striped virtual disk, you must create a new, larger striped virtual disk and then copy the original striped virtual disk's contents to it. You may have to copy the original striped virtual disk's contents to tape first to make room for the new, larger striped virtual disk. This operation must be performed in an offline mode; the data on the virtual disk to be copied cannot be in use.

To change the size of a mirror, first disassemble all images from a mirror before expanding the virtual disk. Refer to the section on managing mirrors for more information. Expand each virtual disk separately, then recreate the mirror. When

you recreate the mirror, you will have to synchronize all secondary images to the primary image. During this operation, you may continue to access the data on the virtual disks forming the mirror, but without the benefits of mirroring until the mirror is reassembled.

To change the size of a cache, first unlink the front–end devices. Then disassemble the cache, which leaves intact the child virtual disk (back–end device) being cached. Expand the virtual disk, as desired, then recreate the cache using the front–end devices that you previously unlinked and the back–end device that you just expanded. During this operation, you may continue to access the data on the back–end device, but without the benefit of caching until the cache is reassembled.

Shrinking a Virtual Disk

To reduce the size of a virtual disk, select the **sysadm** operation `Device -> Disk -> Virtual -> Shrink`. The virtual disk cannot be in use (for example, users accessing a database on the virtual disk) when being shrunk. Use this operation to shrink only partitions and aggregations of partitions. If shrinking a partition, you will be trimming blocks from it. If shrinking an aggregation, you may be trimming blocks from the last partition of the aggregation or you may delete entirely a child partition. If shrinking a virtual disk results in only one unnamed child, the aggregation is removed, leaving only one partition which assumes the aggregation's name.

The number of blocks to be trimmed must be less than the total size of the virtual disk. You cannot shrink striped virtual disks or virtual disks that are images of mirrored disks. Also, you should not shrink a virtual disk that has a file system on it; use the `File system -> Local Filesys -> Shrink` operation instead.

After you select this operation, the system displays the following queries:

`Virtual disk:`
: Select the virtual disk that you want to shrink.

`Blocks to Trim:`
: Enter the number of blocks you want to remove from the virtual disk.

Striped and Mirrored Virtual Disks and Cached Virtual Disks

To change the size of a striped virtual disk, you must create a new, smaller striped virtual disk and then copy the original striped virtual disk's contents to it. You may have to copy the original striped virtual disk's contents to tape first to make room for the new, smaller striped virtual disk. This operation must be performed in an offline mode; the data on the virtual disk to be copied cannot be in use.

To change the size of a mirror, first disassemble all images from a mirror before shrinking the virtual disks that are used as images. Refer to the section on managing mirrors for more information. Shrink each virtual disk separately, then re–create the mirror. When you re–create the mirror, you will have to synchronize all secondary images to the primary image. During this operation, you may continue to access the data on the virtual disks forming the mirror, but without the benefits of mirroring until the mirror is reassembled.

To change the size of a cached virtual, first disassemble the cache before shrinking the virtual disks used as front– and back–end devices. Refer to the section on managing caches for more information. Shrink each cached virtual disk separately, then re–create the cache. During this operation, you may continue to access the data on the back–end device, but without the benefit of caching until the cache is reassembled.

Copying a Readable and Writable Virtual Disk

To copy a virtual disk that is readable and writable, select the **sysadm** operation `Device -> Disk -> Virtual -> Copy`.

NOTE: To copy a read–only virtual disk (such as one on a CD–ROM device), do not use the **sysadm** virtual disk copy operation. Should you attempt to copy such a virtual disk, you will receive an error. Instead, go to the next section for instructions.

The source and destination virtual disks must be the same size. The physical disks on which the source and destination reside must both be registered, and the virtual disks must have different names. The virtual disk's file system may be mounted and in use (for example, users accessing a database on the virtual disk) when being copied. Following a copy operation, two copies exist: the source and the destination. If you copy a mounted virtual disk, the copy will be suitable for mounting as a DG/UX file system.

After you select this operation, the system displays the following queries:

`Source Virtual Disk:`
: Select the virtual disk that you want to copy.

`Destination Virtual Disk:`
: Select the name of the virtual disk you want to copy to. The destination virtual disk must already exist, it must have a name that is different from the source disk, and it cannot be in use.

`Throttle Value in Milliseconds: (0-300) [0]`
: Select a number of milliseconds for the system to wait between successive I/O operations required to complete the copy. Choosing a non–zero value slows down the movement of data and can reduce contention for a device. Larger numbers slow down the copy operation. Choosing zero causes the copy operation to proceed at full speed; thus, slowing down all other users' access to the device.

CAUTION: This operation destroys any data on the destination virtual disk.

Copying a Read–only Virtual Disk

Use this procedure to copy a read–only source virtual disk to a destination virtual disk.

If the source virtual disk is on a CD–ROM or any other read–only medium, you must copy the source virtual disk to a pre–existing virtual disk to be used as the destination.

To create explicitly a destination virtual disk, use the **sysadm** operation `Disk -> Virtual -> Create`, specifying the same attributes for the destination as the source, particularly its size.

Use the **cp** shell command to copy the source virtual disk to the destination virtual disk. Use the following command syntax:

```
# cp /dev/dsk/source-virtual-disk /dev/dsk/destination-virtual-disk
```

where:

cp is the shell command to perform the copy operation. See the **cp**(1) manual page for more information.

/dev/dsk is the standard location for mounted virtual disks. Each time you create a virtual disk, the system automatically logs an entry to this file.

You may perform the copy while the source virtual disk is mounted and in use. The destination virtual disk, however, cannot be mounted. Of course, after you have performed the copy, mount the destination virtual disk and unmount the source virtual disk.

NOTE: Do not use this procedure as a general–purpose operation for copying virtual disks on read–write physical disks. Use the **sysadm** copy virtual disk operation instead.

Moving a Virtual Disk

To move a virtual disk, select the **sysadm** operation `Device -> Disk -> Virtual -> Move`. The source and destination virtual disks must be the same size. The difference between a copy and a move is that for the copy, two copies having different names results. For a move, only one copy results. The source is deleted and its content is copied to the destination, which assumes the name of the source. The virtual disk's file system may be mounted and in use (for example, users accessing a database on the virtual disk) when being moved.

The source and destination virtual disks must be the same size. The physical disks on which they reside must both be registered, and the virtual disks must have different names.

The move operation is useful for salvaging data from a virtual disk on a failing physical disk. You are notified of a failing disk through the report of soft errors to the system console. While the data on the failing device is still in use, you can move its content to another virtual disk without users experiencing interruption in service. The move operation is useful also for reorganizing virtual disks on a physical disk wherein you specify their exact locations on the physical disk.

After you select this operation, the system displays the following queries:

`Source Virtual Disk:`
Select the virtual disk that you want to move.

`Destination Virtual Disk:`
Select the name of the virtual disk you want to move to. The destination virtual disk must already exist, and have a name that is different from the source disk.

`Throttle Value in Milliseconds: (0-300) [0]`
Select a number of milliseconds for the system to wait between successive I/O operations required to complete the move. Choosing a non–zero value slows down the movement of data and can reduce contention for a device. Larger numbers slow down the copy operation. Choosing zero causes the move operation to proceed at full speed; thus, slowing down all other users' access to the device.

CAUTION: This operation destroys any data on the destination virtual disk.

Listing Information about a Virtual Disk

To list (display) information about a virtual disk, select the **sysadm** operation `Device -> Disk -> Virtual -> List`. After you select this operation, the system displays the following queries:

`Virtual Disk(s): [All_visible]`
You can list all named virtual disks, excluding system partitions, which begin with a period (.). Typing ? shows your choices.

`All visible` Presents a report that includes named virtual disks. System partitions (those beginning with a period) are excluded.

`All` Gives a report on all visible disks and unnamed virtual disks, which result from creating an aggregation or cache with new, unnamed child partitions. Unnamed virtual disks are signified by a hexadecimal number that begins with a comma. Also, the report includes system partitions, whose names begin with a period (.).

`Virtual disks` Presents a report for the specific virtual disks named.

`Report type: [one-line]`
The examples in this section illustrate an aggregation named "simpson" with partitions "marge" and "homer." The report types are:

`one-line` presents abbreviated list containing one line per virtual disk. An example follows:

```
Report Type: [one-line]
Name              Volume Temp  Size    Type
homer              V           20000   partition on sd(insc(0),0,0)
marge              V           20000   partition on sd(insc(0),0,0)
simpson            V           40000   aggregation of 2 pieces
sd(insc(0),0,0)               1295922  physical for sd(insc(0),0,0)
```

Under the "Volume" heading, "V" represents a virtual disk that is a volume, "L" means that the virtual disk has only a long–name entry in **/dev/dsk**, which results if multiple volumes have the same name. "T" means that the virtual disk is temporary and will be deleted when the system shuts down. "F" under the "Temp" column means the virtual disk is floating (changes to virtual disk definition will be deleted when the system shuts down.)

`standard` presents paragraph–style report on each virtual disk listed. An example follows:

```
Virtual disk name:  homer
It is usable, in use, a volume, persistent, 20000 blocks.
Its type is:  partition on sd(insc(0),0,0).
It starts at block 621 and is 20000 blocks long
Partitioned virtual disk:
   Virtual disk name:  <no name>
   It is usable, in use, not a volume, persistent, 1295922 blocks.
   Its type is:  physical for sd(insc(0),0,0).

Virtual disk name:  marge
It is usable, in use, a volume, persistent, 20000 blocks.
Its type is:  partition on sd(insc(0),0,0).
It starts at block 20621 and is 20000 blocks long
Partitioned virtual disk:
   Virtual disk name:  <no name>
   It is usable, in use, not a volume, persistent, 1295922 blocks.
   Its type is:  physical for sd(insc(0),0,0).

Virtual disk name:  simpson
It is usable, not in use, a volume, persistent, 40000 blocks.
Its type is:  aggregation of 2 pieces.
Piece 1 of the aggregation:
   Virtual disk name:  homer
   It is usable, in use, a volume, persistent, 20000 blocks.
   Its type is:  partition on sd(insc(0),0,0).
   It starts at block 621 and is 20000 blocks long
Piece 2 of the aggregation:
   Virtual disk name:  marge
   It is usable, in use, a volume, persistent, 20000 blocks.
   Its type is:  partition on sd(insc(0),0,0).
   It starts at block 20621 and is 20000 blocks long

Virtual disk name:  sd(insc(0),0,0)
It is usable, in use, not a volume, persistent, 1295922 blocks.
Its type is:  physical for sd(insc(0),0,0).
```

`recursive` presents information about each selected virtual disk, followed by information about its descendents. An excerpt follows:

```
Name                    Volume Temp  Size Type
simpson                      V      40000 aggregation of 2 pieces
   homer                     V      20000 partition on sd(insc(0),0,0)
      <no name>                    1295922 physical for sd(insc(0),0,0)
   marge                     V      20000 partition on sd(insc(0),0,0)
      <no name>                    1295922 physical for sd(insc(0),0,0)
```

`partitions` reports information about the partitions and free space on each listed virtual disk. Nameless aggregated partitions are designated by the description beneath "Role" in the form: "Piece 1 of n." An excerpt follows:

```
Name                          Size        Free Comments
simpson                      40000       40000 not partitioned

   Partition Name              Role              Address         Size
   <virgin free space>                                 0        40000

Name                          Size        Free Comments
sd(insc(0),0,0)           1295922     1255285

   Partition Name              Role              Address         Size
   <Various System Partitions>                         0          621
   homer                                             621        20000
   marge                                           20621        20000
   <free space>                                    40621      1255285
   <Various System Partitions>                   1295906           16
```

`long` gives an exhaustive, paragraph–style report on each listed virtual disk. Use this report style for striping information. An excerpt follows:

```
Virtual disk name:  homer
It is usable, in use, a volume, persistent, 20000 blocks.
Its type is:  partition on sd(insc(0),0,0).
It starts at block 621 and is 20000 blocks long
Bad block mapping is enabled.
Partitioned virtual disk:
   Virtual disk name:  <no name>
   It is usable, in use, not a volume, persistent, 1295922 blocks.
   Its type is:  physical for sd(insc(0),0,0).
   It is write-enabled.
   It has software bad block mapping facilities.

Virtual disk name:  marge
It is usable, in use, a volume, persistent, 20000 blocks.
Its type is:  partition on sd(insc(0),0,0).
It starts at block 20621 and is 20000 blocks long
Bad block mapping is enabled.
Partitioned virtual disk:
   Virtual disk name:  <no name>
   It is usable, in use, not a volume, persistent, 1295922 blocks.
   Its type is:  physical for sd(insc(0),0,0).
   It is write-enabled.
   It has software bad block mapping facilities.

Virtual disk name:  simpson
```

```
It is usable, not in use, a volume, persistent, 40000 blocks.
Its type is:  aggregation of 2 pieces.
It is striped, stripe size is 16 blocks.
Piece 1 of the aggregation:
   Virtual disk name:  homer
   It is usable, in use, a volume, persistent, 20000 blocks.
   Its type is:  partition on sd(insc(0),0,0).
   It starts at block 621 and is 20000 blocks long
   Bad block mapping is enabled.
Piece 2 of the aggregation:
   Virtual disk name:  marge
   It is usable, in use, a volume, persistent, 20000 blocks.
   Its type is:  partition on sd(insc(0),0,0).
   It starts at block 20621 and is 20000 blocks long
   Bad block mapping is enabled.

Virtual disk name:  sd(insc(0),0,0)
It is usable, in use, not a volume, persistent, 1295922 blocks.
Its type is:  physical for sd(insc(0),0,0).
It is write-enabled.
It has software bad block mapping facilities
```

`parents` The name of the parent virtual disk is listed for the specified child virtual disk. For example, if you specify child partition "marge," the parent listed will be "simpson."

Managing Mirrored Virtual Disks

You can improve the reliability and availability of your Data General AViiON system by mirroring virtual disks. You can set up disk mirroring through the hardware (if you have a disk-array), through the software, or both. This manual covers only mirroring virtual disks through the software. For information on hardware disk mirroring, see your disk-array documentation.

A mirror is composed of up to three virtual disks that are identical images of each other: they all contain the same data. The system manages access to the virtual disks in a manner that is transparent to users. For the administrator, mirrored virtual disks are easy to maintain. Mirrored virtual disks provide three benefits:

Data availability

Because the virtual disks comprising the mirror all contain the same data, a single disk failure no longer interrupts access to the data. As long as one virtual disk in the mirror remains functional, users experience no interruption in service.

Data integrity

With multiple identical images of your data, you greatly reduce the risk of losing data due to hardware failure on one drive.

Performance

Mirrored virtual disks whose images lie on different physical disks offer increased throughput in environments where multiple concurrently running

applications perform intensive reads of the mirror. This benefit arises because the system can use the images of the mirror as individual virtual disks during concurrent read operations, using one image to satisfy one read request while using another image to satisfy a different read request. Thus, the mirror distributes the I/O load across multiple disk drives.

While a single running application will not exhibit increased performance, the system overall will show an improved performance. This benefit does not occur in environments where only one running application reads the mirror at a time, nor does it occur in environments where the images do not reside on different physical disks.

You do not need any special hardware, or software other than the DG/UX system, to take advantage of the benefits of software mirroring; however, mirroring provides the most benefits on systems with multiple physical disks.

Understanding the Concepts

A mirrored virtual disk consists of two or three (maximum) identically sized child virtual disks, called images. The mirror appears as a virtual disk on the system, and can be accessed the same as a virtual disk that is not mirrored. Figure7–1 shows a typical mirrored virtual disk configuration.

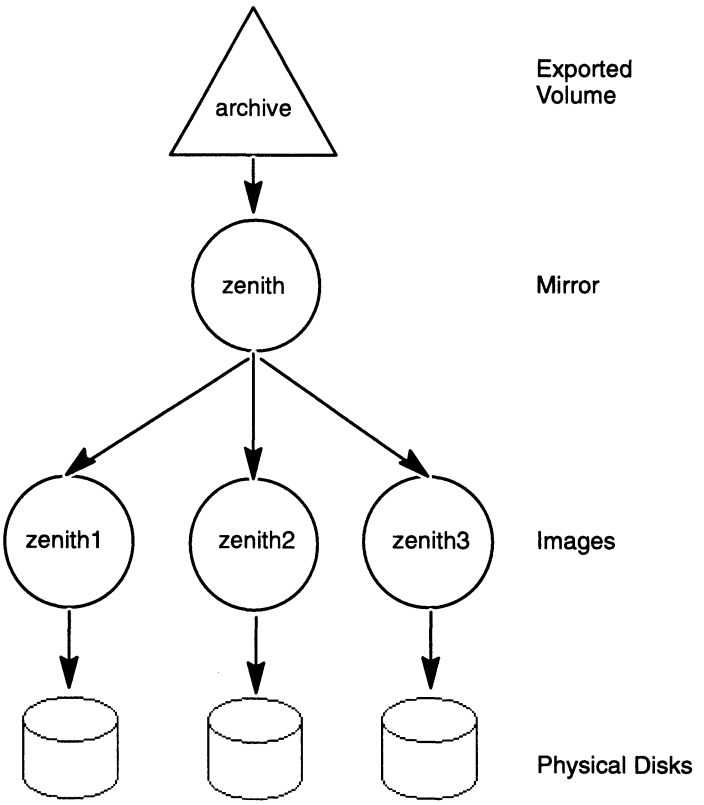

Figure 7–1 Typical Mirrored Virtual Disk Configuration

If the images are named and are volumes, they will also appear in the **/dev/dsk** and **/dev/rdsk** directories, but they may be accessed for read–only operations while they are part of the mirror.

When a user or an application writes to a file to a mirrored virtual disk, the system duplicates the write operation on each mirror image. When a user or application

performs a read operation on a mirrored virtual disk, the system selects one of the images to satisfy the read request.

If a read operation fails on one image, the system satisfies the read request by reading from another image instead. If a bad block caused the original failure, and the physical disk on which the image resides has bad block mapping enabled, the system attempts to repair it by remapping the bad block and updating it with the correct data from a good block on another image. If this repair operation succeeds, the image remains an active member of the mirror. If the repair operation fails, however, the system responds to the failure as it does to a failed write operation, described below.

If a write to any image fails, the system marks the image as corrupt and suspends further use of the image. If there are still functioning images in the mirrored virtual disk, the system will continue to serve read and write requests to the mirror.

When an image becomes corrupt, the system issues a warning. An example follows:

```
Oct 22 10:31:45 homer dg/ux: Warning:  Image
'vdm(marge_image2,2CC7EE0E,0C052A9D,0)' on mirror
'vdm(marge_mirror,2CC7EE30,0C052A9D,0)' has failed (status = 77005171).
```

Probably, a second message will follow when the mirror attempts to update the timestamps of both images, and gets a failure when writing to the failed one. An example of such a message follows:

```
Oct 22 10:31:45 homer dg/ux: Error:  Cannot store the attributes for
virtual disk 'vdm(marge_mirror,2CC7EE30,0C052A9D,0)' on disk
'sd(insc@7(FFF8A000,7),1,0)' (status = 77005171).
```

Once an image becomes corrupt, the system considers it to be *out of sync* with the other images in the mirror. Being out of sync means that the data in the image is inconsistent with the data in the other images: the images are no longer identical. After you fix the problem and restore the image to service, you then need to synchronize the corrupted image with a known good image. The system then copies the master image onto the out-of-sync image.

The system conducts synchronization without interrupting user access to the mirror. The system handles any concurrent user access to the mirror in a fashion that protects data integrity and keeps all images up to date. Except for a possible effect on disk I/O performance, users will not know that synchronization is occurring.

As part of rebooting activities, the system registers physical disks, synchronizes mirror images, and logs a message using the **syslog** error logging facility. The message is from the **kern** facility and is level **warning**. If synchronization fails, the system logs a **kern.err** message, which by default goes to the system console and to **/usr/adm/messages**.

An example of such a warning follows:

```
Oct 22 10:31:45 homer dg/ux: Error:  Cannot store the attributes for
   mirror 'vdm(simpson_mirror,2CE83C55,0C052A9D,0)' from master image
```

```
'vdm(marge_image1,2CE83BAF,0C052A9D,0)' to destination image
'vdm(homer_image2,2CE83BBB,0C052A9D,0)' (status = 4005007).
```

Subsequently, when the mirror attempts to update the timestamp of both images, it gets a failure when writing to the failed one. The following warning is then displayed:

```
Oct 22 10:31:45 psycho dg/ux: Error:  Cannot store the attributes for
   virtual disk 'vdm(simpson_mirror,2CC7EE30,0C052A9D,0)' on disk
   'sd(insc@7(FFF8A000,7),1,0)' (status = 77005171).
```

Booting from a Mirror

To boot an image residing on a mirrored virtual disk, specify the name of one of the available images rather than the name of the mirror itself. For example, if you created a mirrored virtual disk for your **root** virtual disk, and called the images **root1** and **root2**, you would boot by specifying either of the images, as in the following boot command line:

```
SCM> b sd(dgsc(),1)root2:/dgux ↵
```

Operations for Managing Mirrored Virtual Disks

The operations for mirrored virtual disks are available in **sysadm**'s `Device -> Disk -> Virtual -> Mirrors` menu. There are operations to mirror an existing virtual disk, link and unlink a mirror image, unmirror, synchronize mirror images, control the speed of synchronization, modify a mirror, and list disk mirror information.

Considerations for Mirrored Virtual Disks

To mirror an existing virtual disk, select **sysadm**'s `Device -> Disk -> Virtual -> Mirrors -> Mirror` operation. Before you invoke the operation to create a mirror, you need to decide several things:

- How many images will the mirror include?
- On which physical drives and controllers will you put the images?
- How will you configure the virtual disks for each image?
- How many images must be available before mounting the mirror?
- Will the system synchronize images (as needed) at every boot? How quickly should the system perform the synchronization?

How many images?

In deciding how many images will make up the mirror, you need to consider the tradeoff between availability and cost. A mirror comprising three images provides higher availability of data: the mirror can withstand more individual image failures before it becomes completely out of service. On the other hand, you can save disk space by having only two images for the mirror. Consider, for example, that to make a three-image mirror that has the effective size of 300 Mbytes, you need to create three 300-Mbyte virtual disks, making a total resource cost of 900 Mbytes.

Where should you put them?
As you consider which physical drives and controllers will hold your mirror images, you encounter a similar tradeoff between data availability and resource expense. Ideally, you want each image to reside on a different physical disk attached to a different hardware controller. By isolating the images this way, you insulate them from each other's possible hardware failures. Consider, for example, a three-image mirror where all the images are on the same physical disk. If the disk or disk controller fails, the entire mirror becomes inaccessible until you can repair the disk and restore the data from backup. Obviously, placing each image on a different physical disk is the wisest configuration. The same principle applies to the disks' hardware controllers. If the images are all on different disks, but the disks all depend on the same controller, you risk losing the entire mirror if the controller fails.

What kind of virtual disks should you create?
Once you have decided on the number and location of the images, you need to make the decisions that are pertinent any time you create a virtual disk. The only requirement for mirroring is that all virtual disks be the same size. The names of the virtual disks and the number and placement of any virtual disk pieces, however, are completely up to you.

What are the minimum number of images required?
You need to decide how many images must be available before you can mount the mirror and make it accessible to users. This question is just another way of asking how many corrupt images you will tolerate.

The data in a mirror image can be in either of three states:

- completely good
- in the process of being synchronized and will eventually be good
- corrupt

If your mirror has three images and you require at least two functional images, this means that the data is either completely good or is in the process of being synchronized. If you are more interested in availability than reliability, with one corrupt image, one being synchronized, only one good image, you will be using only one image for your transactions. Should you trade off reliability for availability, you may decide to delay your transaction processing until the image has completed synchronization. You will then have two truly good, functional images.

Keep in mind that you may not be on hand to evaluate the state of the system after a system crash. After a power outage, for example, a Data General AViiON system can reboot itself without operator intervention as soon as power returns. Depending on how you configured the system, it may come up all the way to run level three, where local users and OS clients may begin work and access file systems. The availability requirement you set will determine whether or not users can access a mirror that may have lost an image in the crash. If the data in the mirror is such that you do not want it in use if there is only one functioning image, you should set the availability requirement to two or three.

Setting the availability limit allows you to enforce either a policy of high data availability, where downtime is an issue, or a policy of high data integrity, where you depend on redundant images to protect from data loss. For high availability, select a lower integrity requirement. For high integrity, select a higher availability requirement.

Note that when counting "in-sync" images, any image that is the destination of an automatically started synchronization operation (see later section on automatic synchronization) is considered to be in-sync for the purposes of deciding if a sufficient number of images is available to put the mirror into use. For example, consider a mirror with three images—one is correct and two are out-of-sync—for which auto-sync is enabled. In this case, two synchronize operations are started, and the mirror is immediately considered to have three noncorrupt images. You cannot tell the system to wait for the automatic synchronize operations to complete before placing the mirror into service.

How many lost images will be tolerated?
The previous selection—the minimum number of images required—might be considered the high water mark, whereas the maximum number of images lost might be thought of as the low–water mark. You are balancing the minimum number of images you will operate with and the maximum number of images you are willing to operate without. An image is considered "lost" if it can't be located. An image may be lost if its host disk is disconnected or if the disk is unavailable because of a crash, for example.

The value you select for this parameter is again a question of reliability versus availability. Suppose you tolerate one lost image, which you know to contain the most up–to–date data. If you tolerate this loss, the remaining image(s) will still be available for use, but the data integrity may be questionable. Alternatively, if you choose reliability as your aim, you must halt service until the damaged disk is back on–line, sacrificing availability. To maximize reliability, you may want to choose 0 as the maximum number of lost images to tolerate; for availability, perhaps, 1. Of course, this selection must be made with consideration of the minimum number of images required.

Should the system perform synchronization automatically?
When the system boots, it examines the images of each mirror that it finds. If it appears that the images may be out-of-sync with each other (for example, due to the system crashing), the system can be configured to automatically initiate a synchronize operation to make the images identical.

You set the synchronization speed when you create the mirror. You can change the speed later through the modify operation.

Outline for Creating a Software Mirror

The outline for creating a software mirror follows.

1. Create, or have available, up to three (maximum) identically sized virtual disks.

   One virtual disk will be the primary image; the others (one or two) will be the secondary mirror images. For reliability, each virtual disk should reside on

different controllers and physical disks. The primary image may already contain a file system and data. Although, if you are just creating the virtual disk to be used as the primary image, make sure you put a file system on it. Do not put file systems on the secondary images.

2. Create the mirror.

   Identify the virtual disk to serve as the primary image, and name the mirror. By convention, the primary image's name is assigned to the mirror, which leaves the primary image nameless. The primary image must then be renamed. You are not constrained, however, to using this naming scheme. You can name the mirror and the primary image anything you wish as long as they're unique. You will also supply various mirror attributes.

3. Identify the virtual disks to be linked as secondary images to the mirror.

4. Synchronize the secondary images with the primary image; the contents of the primary image are copied to the secondary images. The primary image for synchronization should be the virtual disk that has the file system. You may synchronize (copy the primary image to a secondary image) after each image is linked. Otherwise, delay synchronizing until all images have been linked.

5. Mount the mirror's file system.

You may now use the mirror as a normal file system.

Once you add virtual disks to a mirror, they are no longer accessible on your system as individual virtual disks. Instead, the mirror exists as a single virtual disk representing them. Any operations that you may perform on a virtual disk, you may perform on the mirror. File system operations such as **dump2** and **restore** function on the mirror's file system the same way that they function on any other file system.

Mirroring a Virtual Disk

You must have a pre–existing virtual disk to mirror and pre–existing virtual disks to use as mirror images. The virtual disks used as the primary image and the secondary mirror images must be the same size.

To build a mirror from existing virtual disks, select the **sysadm** operation `Device -> Disk -> Virtual -> Mirrors -> Mirror`.

After you select this operation, the system displays the following queries:

`Virtual Disk`
: Select the virtual disk whose data you want to mirror.

`Minimum Images Required: (1-3) [1]`
: Enter the number of mirror images that must be available and "in sync" or in the process of being synchronized before the operating system makes the software mirror available for mounting or opening. If the value is 1, the mirror will be made available as soon as the first mirror image is present. The higher the value, the greater the data reliability, which is at the possible expense of delays in making the mirror available.

`Maximum Number of Lost Images to Tolerate: (0-3) [2]`

Select the number of "lost" mirror images (which reside on physical disks that are not registered) to be tolerated during normal system operation. At reboot, the operating system calculates how many images a mirror had when it was last in use. The system does not attempt to resynchronize yet. If the number of missing images exceeds the value you select here, the system will not resynchronize the images and, therefore, will not make the mirror available for use. Choosing 0 means that no missing images will be tolerated. If, for example, you select 0 as the number of lost images to tolerate, and if an image is determined to be lost, the mirror will be made unavailable for use. If, as another example, you tolerate the absence of two images (the default), you will operate with only one image available. Consider the tradeoff between availability and reliability when making your decision.

`Automatically Synchronize on System Boot? [yes]`

When the system is brought up, the mirror images may be in synchronization or they may be out of synchronization. If they are out of sync, the operating system can automatically initiate a synchronize operation, so that the mirror will become available for use without your intervention. Choose whether or not you want the operating system to automatically initiate synchronization at boot time.

`Throttle Value in Milliseconds: (0-300) [0]`

Select a number of milliseconds for the system to wait between successive I/O operations required to complete the synchronization. Choosing a non–zero value slows down the movement of data and can reduce contention for a device. Larger numbers slow down the synchronize operation. Choosing zero causes the synchronize operation to proceed at full speed; thus, slowing down all other users' access to the device. This value is used with synchronizations that the system starts automatically, and is also used as a default for manually started synchronizations. You can change this value later through the modify operation.

`New Name for Mirror: [journal]`

Specify the name for the new mirror. Normally, the mirror assumes the name of the virtual disk that is being mirrored (for example, **journal**), and the virtual disk being mirrored is given a new name (next prompt). However, you may give the mirror and the image any names you please, as long as they are unique.

`New Name for Image:`

Specify the name for the virtual disk image. You are renaming the virtual disk that is being mirrored. Typically, the name of the virtual disk being mirrored travels with the mirror, and you must select a new name to associate with the mirror image. For example, the name **journal** traveled with the mirror, so you might assign a new name, **entries–1**, to the virtual disk that is now serving as an image. However, you may give the mirror and the image any names you please, as long as they are unique.

`Do you want to specify another image for this mirror? [no]`
: At this point, you have a one–image mirror. To create another image. answer "yes." Otherwise, answer "no."

`Child Virtual Disk:`
: Select an existing, identically sized virtual disk to serve as a secondary image.

`Begin Sync Immediately? [yes]`
: Specify whether or not you want to immediately synchronize the secondary image with the primary image. If you do not intend to add another secondary image (next prompt), answer "yes." Otherwise, answer "no."

NOTE: Later, you may select explicitly an operation to synchronize the desired images. A later section discusses this operation.

`Do you want to specify another image for this mirror? [no]`
: Specify whether or not you want to add another secondary image to the mirror. The maximum number of images per mirror is three. If you answer "yes," the preceding two prompts repeat. Otherwise, no more prompts are delivered, and the new mirror is built according to your selections.

Linking One or More Images to an Existing Mirror

To link one or more images to an existing mirror, select the sysadm operation `Device -> Disk -> Virtual -> Mirrors -> Link`. The image (child) virtual disk to be linked must be the same size as the existing mirror images.

After you select this operation, the system displays the following queries:

`Existing Mirror:`
: Select the mirror to link another image to.

`Child Virtual Disk:`
: Select an existing virtual disk to serve as another image to link to the mirror.

`Begin Sync Immediately? [yes]`
: Specify whether or not you want to immediately synchronize the newly linked secondary image with the primary image. If you intend to link yet another image to the mirror, rather than synchronizing the newly linked image now, you may defer synchronization until after you have linked the next image. Then you may synchronize both images at once, which will result in better performance. Answer "no" to postpone synchronization. Go to the final prompt in this section.

  Later, you may select explicitly an operation to synchronize the newly linked secondary image with the mirror. A later section discusses this operation.

  Otherwise, if you do not wish to link another secondary image, answer "yes"to begin synchronization immediately and the next prompt appears.

`Throttle Value in Milliseconds: (0-300) [0]`

Select a number of milliseconds for the system to wait between successive I/O operations required to complete the synchronization. Choosing a non–zero value slows down the movement of data and can reduce contention for a device. Larger numbers slow down the synchronize operation. Choosing zero causes the synchronize operation to proceed at full speed; thus, slowing down all other users' access to the device.

`Do you want to specify another image for this mirror? [no]`

To create another secondary image, answer "yes." **Sysadm** repeats all the prompts again to collect information about the next image. Recall that a mirror is restricted to three images. Otherwise, if you have finished linking secondary images, accept the default "no."

The desired image(s) are linked to the mirror and synchronized, as appropriate.

Unlinking an Image from a Mirror

To unlink (remove) an image from an existing mirror, select the **sysadm** operation `Device -> Disk -> Virtual -> Mirrors -> Unlink`. Remember that when you unlink an image, you must be left with the minimum image(s) required. For example, if you originally set up a three–image mirror to contain a minimum of one image, unlinking one image will leave two functional images. Unlinking two images, however, violates the minimum number of images required.

If the image to be unlinked contains a file system, the operating system will flush all modified buffers to all mirror images before the unlink operation is performed. Consequently, you can mount separately and use an unlinked mirror with assurance that the data is up–to–take.

After you select this operation, the system displays the following queries:

`Existing Mirror:`

Select the mirror from which to unlink the image.

`Image Virtual Disk(s) to Unlink:`

Select an existing virtual disk image to unlink from the mirror. Remember that at least one synchronized image must remain in the mirror

Following your confirmation, the desired image is unlinked (removed) from the mirror. The virtual disk still exists, but is not associated with the mirror.

Mirroring Critical Data for an Online Backup: A Typical Scenario

The following example shows a typical **sysadm** session for mirroring a virtual disk that contains critical data to be backed up on a regular basis. You have a striped aggregation named **zenith** that you want to duplicate in a two–image mirror. When it's time to perform an online backup, you can then unlink one of the images and dump it to tape while the mirror is online and in use.

To build a mirror from an existing virtual disk, select the **sysadm** operation `Device -> Disk -> Virtual -> Mirrors -> Mirror`.

```
Virtual Disk: zenith ↵
Minimum Images Required: (1-3) [1] ↵
Maximum Number of Lost Images to Tolerate: (0-3) [2]↵
Automatically Synchronize on System Boot? [yes] ↵
Throttle Value in Milliseconds: (0-300) [0]↵
New Name for Mirror: [zenith]↵
New Name for Image: zenith1↵
Do you want to specify another image for this mirror? [no] yes ↵
Child Virtual Disk: zenith2
Begin Sync Immediately [yes]
Do you want to specify another image for this mirror? [no] ↵
Virtual disk "zenith,2CBC65C1,0C052A9D,0" renamed to "zenith1"
Virtual disk inserted at zenith.
Child virtual disk made a volume.
Virtual disk "zenith2" linked as a child to virtual disk "zenith".
Synchronization started on virtual disk disk mirror "zenith".
```

Next, unlink image **zenith2** to dump to tape. The remaining mirror image remains online and in use while one image is unlinked.

To unlink an image from a mirror, select the **sysadm** operation `Device -> Disk -> Virtual -> Mirrors -> Unlink`.

```
Existing Mirror: zenith↵
Image Virtual Disk(s) to Unlink: zenith2
Child virtual disk "zenith2" unlinked from parent virtual disk "zenith"
```

After you mount **zenith2**'s file system, you can then use the desired tool such as the Legato Networker™ product or the **tar**, **cpio**, or **dump2** commands to back up the contents of the file system to tape. Following the backup, you may want to relink the image to the mirror and synchronize the data, keeping the mirror image available for the next backup.

Unmirroring (Dismantling) a Mirror

To dismantle all images from a mirror, bar one, select the **sysadm** operation `Device -> Disk -> Mirrors -> Mirror -> Unmirror`. Unmirroring leaves behind one child virtual disk image, which assumes the name of the mirror. You may disassemble the mirror while it is in use, without interruption to access of the data, but you cannot unmirror a disk mirror while synchronization is in progress.

This operation deletes the mirror and passes on the mirror name to a mirror image that you specify. A child image's inheritance of its parent mirror's name prevents a need to remount the child image's file system. If the child did not inherit its parent's name, removal of the mirror does require that you unmount the mirror's file system and mount the child image's file system in the mirror's place. No unmounting and mounting is required if a child assumes the parent's name.

After you select this operation, the system displays the following queries:

`Existing Mirror:`
: Select the mirrored virtual disk to unmirror.

`Image Virtual Disk to Get Name [`*mirror-name*`]:`
Select the child image that will inherit the parent mirror's name. Applications that are currently accessing the data on this image will use this renamed image.

This operation removes the mirror, unlinks the images from the mirror, and assigns the parent's name to a child image. You may remove remaining child images.

After you have disassembled a mirror, you can use its virtual disks just as you would any virtual disks.

Synchronizing a Mirror's Images

Under normal circumstances, the system maintains the exact same data in all images in a mirror. Under some conditions, however, one image in a mirror may become inconsistent with the rest. An inconsistent image is considered "out of sync" or "corrupt." When this happens, use the Synchronize operation to bring the inconsistent image up to date. When an image is out of sync, the operating system will not use that mirror's image. Or, if the required number of synchronized images cannot be met, the operating system will make the entire mirror inaccessible. In this case, you must resynchronize the mirror's images. You do not need to unmount a mirror to synchronize an image in it.

You can synchronize multiple mirror images at once but from only one primary image. It is more efficient to synchronize multiple images at once rather than separately.

To synchronize one or more images, select the **sysadm** operation `Device -> Disk -> Virtual -> Mirrors -> Synchronize`. The image (child) virtual disk to be linked must be the same size as the existing mirror images.

After you select this operation, the system displays the following queries:

`Existing Mirror:`
Select the mirror whose image(s) need to be synchronized.

`Throttle Value in Milliseconds: (0-300) [0]`
Select a number of milliseconds for the system to wait between successive I/O operations required to complete the synchronization. Choosing a non–zero value slows down the movement of data and can reduce contention for a device. Larger numbers slow down the synchronize operation. Choosing zero causes the synchronize operation to proceed at full speed, slowing down all other users' access to the device.

`Master Image Virtual Disk:`
Select the image that will serve as the master, or primary, image for copying. The primary image is the one presumed to be up–to–date and correct.

While synchronization is in progress, the system continues to allow access to the mirror. You cannot break a mirror on which synchronization is occurring.

`Sync Master Image To: [All out-of-sync images]`
Allow the operating system to determine the images that need to be synchronized or specify explicitly the name(s) of images to be synchronized.

CAUTION: Be very careful when selecting the master, or primary, and secondary images for the synchronization operation. If you accidentally specify a good image as the secondary, the operation will destroy all data on the image.

The system then copies the entire master image block-by-block onto the secondary image(s), overwriting it entirely. The system displays a warning that the procedure will destroy the contents of the image(s) being synchronized. If you decide against resynchronization, answer "no" to continue. Otherwise, answer "yes," and the operation begins.

At the beginning and end of a synchronization session, the system logs a message using the **syslog** error logging facility. The message is from the **kern** facility and is level **warning**. By default, **kern.notice** messages go to the system console and to the file **/usr/adm/messages**. If synchronization fails, the system logs a **kern.err** message, which by default goes to the system console and to **/usr/adm/messages**.

An example of such a warning follows:

```
Oct 22 10:31:45 psycho dg/ux: Warning:  Image
'vdm(gth_image2,2CC7EE0E,0C052A9D,0)' on mirror
'vdm(gth_mirror,2CC7EE30,0C052A9D,0)' has failed (status = 77005171).
```

Subsequently, when the mirror attempts to update the timestamp of both images, it gets a failure when writing to the failed one. The following warning is then displayed:

```
Oct 22 10:31:45 psycho dg/ux: Error:  Cannot store the attributes for
   virtual disk 'vdm(gth_mirror,2CC7EE30,0C052A9D,0)' on disk
   'sd(insc@7(FFF8A000,7),1,0)' (status = 77005171).
```

To change how the system handles **warning** messages, see Chapter 16 and the manual page for **syslog.conf**(5).

Adjusting Synchronization Speed (Throttling)

The throttle synchronization operation allows you to adjust the speed of a synchronization in progress by specifying a throttle value.

A throttle value is an amount of time (in milliseconds) that separates each I/O operation that comprises the synchronization operation. A throttle value of 0 means that the sync will be performed as fast as possible (I/O operations are continuous), which can consume significant system resources. A non–zero throttle value will cause the synchronization to take longer (an amount of time separates each I/O operation), but the system as a whole will be slowed less. Previous revisions of DG/UX offered a "slow" speed for mirror synchronizations, which corresponds to a throttle value of approximately 100.

To change the throttle value for a synchronization operation that is in progress, select the **sysadm** operation `Device -> Disk -> Virtual -> Mirrors -> Throttle Sync`.

After you select this operation, the system displays the following queries:

`Throttle Value in Milliseconds: (0-300) [0]`

Select a number of milliseconds for the system to wait between successive I/O operations required to complete the synchronization. Choosing a non–zero value slows down the movement of data and can reduce contention for a device. Larger numbers slow down the synchronize operation. Choosing zero causes the synchronize operation to proceed at full speed; thus, slowing down all other users' access to the device.

`Ongoing Sync: [`*image-name*`]`

Specify the name of the image whose synchronization speed you wish to adjust.

After you confirm your desire to change the synchronization speed, the system performs the operation and displays a message, such as the following:

```
Synchronization of image virtual disk "foobar-1" throttled.
```

Halting a Synchronization in Progress

To halt a synchronization that is in progress, select the **sysadm** operation `Device -> Disk -> Virtual -> Mirrors -> Halt Sync`. The image whose synchronization you halt is unusable.

After you select this operation, the system displays the following query:

`Ongoing Sync: [`*image-name*`]`

Specify the name of the image whose synchronization you wish to terminate.

The system performs the operation and displays a message, such as the following:

```
Synchronization of image virtual disk "foobar-1"
terminated.
```

Modifying a Mirror's Attributes

To alter any of these mirror attributes: minimum images required, maximum number of lost images to tolerate, synchronization upon reboot, and default throttle value, select the **sysadm** operation `Device -> Disk -> Virtual -> Mirrors -> Modify`.

After you select this operation, the system displays the same prompts that are delivered when mirroring a virtual disk.

```
Existing Mirror:

Minimum Images Required: (1-3) [1]

Maximum Number of Lost Images to Tolerate: (0-3) [0]

Automatically Synchronize on System Boot? [yes]
```

```
Throttle Value in Milliseconds: (0-300) [0]
```

Refer to the preceding section on mirroring virtual disks for explanations of each prompt.

Listing Software-Mirrored Disk Information

To list information about disk mirrors and their images, select the **sysadm** operation `Device -> Disk -> Virtual -> Mirrors -> Mirror -> List`.

After you select this operation, the system displays the following queries:

`Mirrored Virtual Disks to List: [all]`
: Select all virtual disks or specify by name those you wish to see.

`List Images? [no]`
: Answering "no" lists only the mirror images; "yes," the mirrors' constituent images as well.

When you list mirrors only, the information displayed largely corresponds to the attributes you gave the mirror when you created it. Besides the name, the "Volume" value specifies whether or not it has been made into a volume: "V" represents a virtual disk that is a volume, and "L" means that the virtual disk has only a long–name entry in **/dev/dsk**. "T" means that the virtual disk is temporary and will be lost on system shutdown. "F" means the virtual disk is floating (changes to virtual disk definition will be lost on system shutdown.). It also includes the size of the mirror (that is, the usable size, not the cumulative size of the component images), and the number of images present

A typical simple mirror listing follows:

```
Name                        Volume Temp  Size Type
wrightbros                      V      10000 mirror with 2 images (1 good)
sam                             V       2100 mirror with 2 images (2 good)
```

The display also includes the current status of the mirror and its images. The current status information includes the sync status of any images: whether the image is synchronized or unsynchronized (that is, corrupt) and, if synchronization is in progress, the percentage of synchronization complete and the name of the image acting as the sync master.

A typical full (mirror and images) listing follows:

```
Virtual disk name:  wrightbros
It is usable, not in use, a volume, persistent, 10000 blocks.
Its type is:  mirror with 2 images (1 good).
Minimum required image count is 1, maximum lost image count is 0.
Auto-sync on system boot is enabled, throttle is 0 msec.
Image 1 of the mirror:
   It is in sync.
   Virtual disk name:  orville
   It is usable, in use, a volume, persistent, 10000 blocks.
   Its type is:  partition on sd(insc(0),0,0).
   It starts at block 24447 and is 10000 blocks long
```

```
   Bad block mapping is enabled.
Image 2 of the mirror:
   It is out of sync.
   Virtual disk name: wilbur
   It is usable, in use, a volume, persistent, 10000 blocks.
   Its type is:  partition on sd(insc(0),0,0).
   It starts at block 34447 and is 10000 blocks long
   Bad block mapping is enabled.
```

Creating Software-Mirrored System Disks

DG/UX 5.4R3.00 allows you to create software mirrors for system virtual disks (typically **root**, **usr**, and **swap**) while they are online and in use. Unlike previous revisions of the DG/UX system, you can use the stand–among disk management facility (in this revision, stand–among **sysadm**) instead of the stand–alone disk management utility to create software mirrors for the system virtual disks.

You want to create another set of virtual disks for the system virtual disks **root**, **usr**, and **swap** on separate physical disks. The following table identifies the goals for the mirrored virtual disks.

Mirror name	Image 1	Size in blocks	Image 2	Size in blocks
root	**root1**	40,000	**root2**	40,000
usr	**usr1**	240,000	**usr2**	240,000
swap	**swap1**	50,000	**swap2**	50,000

A typical **sysadm** session to create a mirror for the system virtual disks follows.

Before you start, realize that the system virtual disks are named **root**, **usr**, and **swap**. When you create the mirror for each system virtual disk, the mirror will inherit the name of its first image, and the first image must take on a new name. For example, when mirroring the **root** virtual disk, the mirror will inherit the virtual disk's name, **root**, and the first image must assume a new name, such as **root1**. Naturally, the second image of the **root** virtual disk mirror would be **root2**. The mirror itself must assume the name of the virtual disk being mirrored.

To create a second virtual disk, **root2**, to mirror **root**, perform the **sysadm** operation `Device -> Disk -> Virtual -> Create`.

```
New Virtual Disk Name: root2 ↵
Striped? [no] ↵
Create File System? [yes] no ↵
Select Space by: [Size alone] Disk to partition and partition size ↵
Disk to Partition From: sd(insc(0),0,0) ↵
Length of Piece in Blocks: (1-1295279) 40000 ↵
Starting Block (optional): ↵
Virtual disk "root2" created.
Virtual disk "root2" made a volume.
```

To create a second virtual disk, **usr2**, to mirror **usr**, perform the **sysadm** operation `Device -> Disk -> Virtual -> Create`.

```
New Virtual Disk Name: usr2 ↵
Striped? [no] ↵
Create File System? [yes] no ↵
Select Space by: [Size alone] Disk to partition and partition size ↵
Disk to Partition From: sd(insc(0),0,0) ↵
Length of Piece in Blocks: (1-1295279) 240000 ↵
Starting Block (optional): ↵
Virtual disk "usr2" created.
Virtual disk "usr2" made a volume.
```

To create a second virtual disk, **swap2**, to mirror **swap**, perform the **sysadm** `Device -> Disk -> Virtual -> Create` operation.

```
New Virtual Disk Name: swap2 ↵
Striped? [no] ↵
Create File System? [yes] no ↵
Select Space by: [Size alone] Disk to partition and partition size ↵
Disk to Partition From: sd(insc(0),0,0) ↵
Length of Piece in Blocks: (1-1295279) 50000 ↵
Starting Block (optional): ↵
Virtual disk "swap2" created.
Virtual disk "swap2" made a volume.
```

You now have a duplicate set of empty virtual disks to be used as secondary images for the system virtual disks. Next, you must create three separate mirrors, specifying the mirror name along with the primary and secondary images. To create a mirror, using **root** and **root2** as images, perform the **sysadm** operation: `Device -> Disk -> Virtual -> Mirrors -> Mirror` operation.

```
Virtual Disk: root ↵
Minimum Images Required: (1-3) [1] 2 ↵
Maximum Number of Lost Images to Tolerate: (0-3) [0]↵
Automatically Synchronize on System Boot? [yes]↵
Throttle Value in Milliseconds: (0-300) [0]↵
New Name for Mirror: [root]↵
New Name for Image: root1↵
Another image? [yes]↵
Child Virtual Disk: root2 ↵
Do you want to specify another image for this mirror? [no]↵
Virtual disk "root,2CBC3856,0C052A9D,0" renamed to "root1"
Virtual disk inserted at root.
Child virtual disk made a volume.
Virtual disk "root2" linked as a child to virtual disk "root".
Synchronization started on virtual disk disk mirror "root".
```

To create a mirror, using **usr** and **usr2** as images, perform the **sysadm** operation: `Device -> Disk -> Virtual -> Mirrors -> Mirror` operation.

```
Virtual Disk: usr ↵
Minimum Images Required: (1-3) [1] 2↵
Maximum Number of Lost Images to Tolerate: (0-3) [0]↵
Automatically Synchronize on System Boot? [yes]↵
Throttle Value in Milliseconds: (0-300) [0]↵
New Name for Mirror: [usr]↵
New Name for Image: usr1↵
Do you want to specify another image for this mirror? [no] yes↵
Child Virtual Disk: usr2 ↵
Do you want to specify another image for this mirror? [no] ↵
Virtual disk "usr,2CBC3856,0C052A9D,0" renamed to "usr1"
Virtual disk inserted at usr.
Child virtual disk made a volume.
Virtual disk "usr2" linked as a child to virtual disk "usr".
Synchronization started on virtual disk disk mirror "usr".
```

To create a mirror, using **swap** and **swap2** as images, perform the **sysadm** operation: `Device -> Disk -> Virtual -> Mirrors -> Mirror` **operation.**

```
Virtual Disk: swap↵
Minimum Images Required: (1-3) [1] 2↵
Maximum Number of Lost Images to Tolerate: (0-3) [0]↵
Automatically Synchronize on System Boot? [yes]↵
Throttle Value in Milliseconds: (0-300) [0]↵
New Name for Mirror: [swap]↵
New Name for Image: swap1↵
Do you want to specify another image for this mirror? [no] yes↵
Child Virtual Disk: swap2↵
Do you want to specify another image for this mirror? [no]↵
Virtual disk "swap,2CBC3856,0C052A9D,0" renamed to "swap1"
Virtual disk inserted at swap.
Child virtual disk made a volume.
Virtual disk "swap2" linked as a child to virtual disk "swap".
Synchronization started on virtual disk disk mirror "swap".
```

Booting from a Mirrored Root

To boot your system from a mirrored **root** virtual disk, specify the name of one of the images rather than the name of the mirror itself because the bootstrap can boot from only a partition or an aggregation, not a mirror. For example:

```
SCM> b sd(cisc(),0,0)root1:/dgux ↵
```

Managing Cached Virtual Disks

Disk caching associates two virtual disks, typically one on a small, fast device (front end) with another on a large, slow device (back end), so that an application uses the fast device for read and write operations while the operating system duplicates these operations on the larger device. The purpose of the configuration is to accelerate file system access for I/O-intensive applications without risking data integrity. Figure 7–2 shows a typical caching configuration.

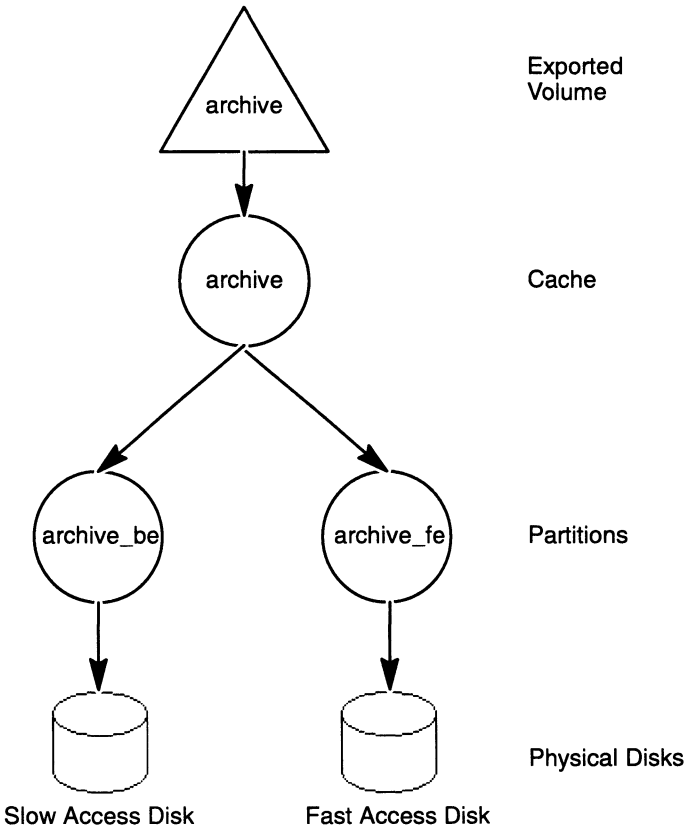

Figure 7–2 Typing Caching Configuration

The primary caching configuration uses a nonvolatile random access memory (NVRAM) or a battery backed-up random access memory (BBURAM) board functioning as the fast and stable front–end device and a physical disk functioning as the slow back–end device. Although the RAM board has a relatively small storage capacity, its superior I/O performance can boost the performance of I/O-intensive applications such as database management systems. In addition, DG/UX 5.4R3.00 supports the use of a disk drive, preferably a fast one, as a front–end device.

These flavors of caching configurations are supported:

- 1:1 correspondence between a single back–end device and a single front–end device.
- *x*:1 correspondence between multiple back–end devices and a single front–end device, where *x* back–end devices can share only one front–end device.

- *x*:*x* correspondence between multiple back–end devices and multiple front–end devices. Multiple front–end devices on a single NVRAM board and/or a fast disk may be linked with the desired number of back–end devices.

Not only is a 1:1 back end–to–front end configuration supported, but also a *x*:1, where *x* represents multiple back ends to one front end, depending on the size of the NVRAM board. If you purchased a NVRAM board for caching, you may optimize its use as a front–end device by associating multiple back–end devices with it. Also, depending on its capacity, you may create multiple virtual disks to be used as front–end devices on the NVRAM board or fast disk to be associated with the desired back ends. Figure 7–3 shows the sharing of one front–end device with two back ends.

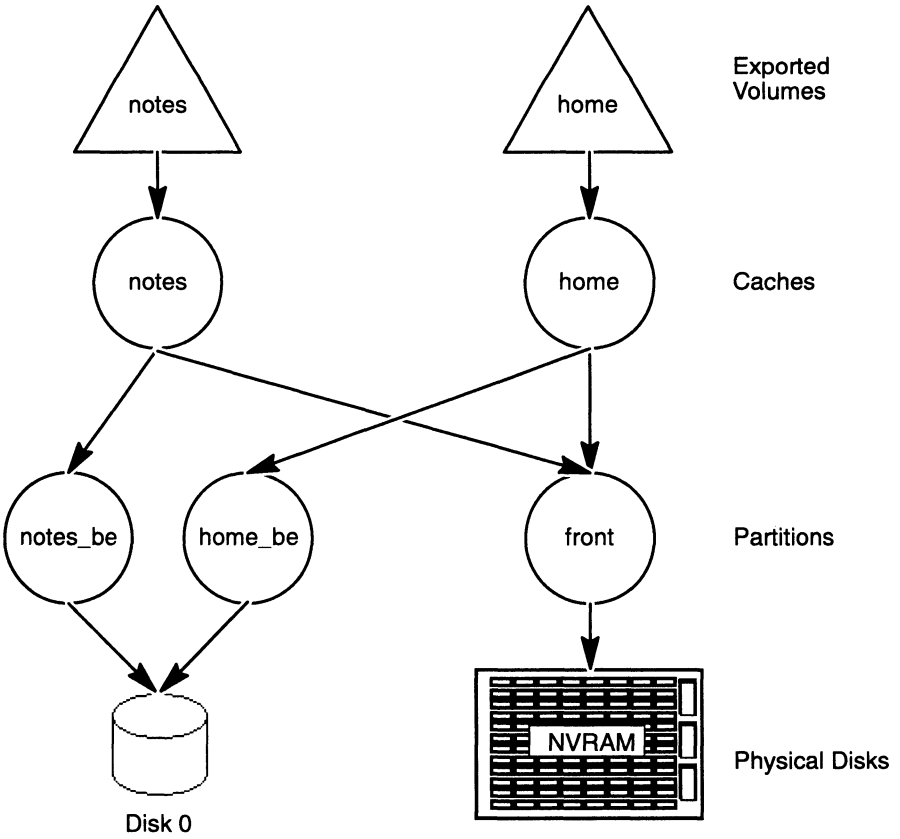

Figure 7–3 Cache Sharing of One Front End with Two Back Ends

RAM-based caches introduce the risk that a failure could lose the data in the cache before the system has a chance to write it to the more stable back end device. The ideal front end device is nonvolatile or battery backed-up RAM or a fast disk, which provides the required speed as well as stability.

The virtual disk functioning as the back end, meanwhile, provides greater storage capacity than a RAM device or disk device and has the added stability normally attributed to disk drives. The cache and its front– and back–end devices must be on local devices; however, they can be accessed by other remote systems through ONC/NFS.

The DG/UX system optimizes disk caching for accessing DG/UX file systems rather than for other data structures (such as databases built directly on virtual or physical

disks). You may, nevertheless, use cached virtual disks for any purpose that benefits from the accelerated I/O performance. You may consider running your applications both with and without disk caching, then comparing results to see which configuration offers the best performance. The following section tells how to get the most out of a cached disk configuration.

How Caching Works

As I/O requests arrive for the cached virtual disk, the system allocates buffers in the front–end device to hold data for the back end device, or disk. These allocated buffers are considered either *clean* or *dirty*. A clean buffer is one whose data matches the corresponding buffer on disk. For example, a buffer that was copied from disk to cache for a read operation is considered clean because it contains the same data that is on the disk. A dirty buffer contains data that is inconsistent with the disk. For example, a buffer that was written by an application but has not yet been flushed to disk is considered dirty.

The process of freeing buffers involves seeking out which buffers are the least frequently accessed and flushing their contents to disk (if dirty) and then flagging them as unallocated. The system is then free to allocate them for more I/O requests.

To determine which buffers are the least frequently accessed, the system maintains a single *weight number* for each buffer. Each time either a read or write I/O request accesses the buffer, the system increments its weight number. When the time comes to reclaim buffers, the system can then compare the weight numbers of the buffers to see which are most frequently accessed (and should stay in the cache) and which are least frequently accessed (and may be freed).

When you shut down the system (using the **shutdown** and **halt** commands), access to the cached disk stops just as for any other disk. The cache then flushes its dirty buffers to disk, after which it flags all buffers as clean.

If an abnormal failure, such as a crash or panic, causes the system to stop without executing the normal **halt** sequence, the cache may be left containing dirty buffers that were never flushed to disk. In this case, the cache device maintains the data until the DG/UX system reboots. After the system reboots, the buffers are recovered.

Building a Cached Virtual Disk

Building a cache involves linking a pre–existing virtual disk (back–end) with a new virtual disk to serve as the front–end. You create the front–end just as you do a virtual disk: by partitioning an existing virtual disk or by using entirely an existing virtual disk. The cache itself is an abstraction. You refer to the entire configuration as a cache; but, the action really occurs between the front– and back–end devices.

The front–end device must be configured already in the kernel. The device driver name for the NVRAM board is **nvrd()**.

The cache itself inherits its name from the back–end device, which leaves the back end nameless. You may rename the back end, or you may leave it nameless. A

back–end device name is important only when you need to refer to it specifically, when modifying or unlinking, for example. Then, of course, the virtual disk created for use as a front end must be named.

If you intend to create a cache in which one front end if connected to one back end, by convention, you may name the cache **foo**; the back end—**foo–be**; and the front end—**foo–fe**. If, however, you intend to create multiple caches to share the same front–end, you may adopt a more generic naming convention for the front end. For example, the first cache and back end might be named **foo** and **foo–be**; the second cache and back end might be named **bar** and **bar–be**; and the front end for both might be named **front**.

Besides identifying the back and front ends, you must also tune several parameters to optimize cache efficiency. Altogether, these tuned parameters form a cache policy.

Caching a Virtual Disk

To cache an existing virtual disk, perform the **sysadm** operation: `Device -> Disk -> Virtual -> Caches -> Cache` operation.

After you select this operation, the system displays the following query:

`Virtual Disk:`
: Identify the name of the pre–existing virtual disk on the slow back–end device that you want to cache.

`Cache Reads? [yes]`
: Specify whether or not to enable reading from the cache. Answering "yes" causes data to be shipped from the back–end device to the cache for reading. Otherwise, data is read directly from the slow back–end device, which is referred to as pass–through mode for reading.

`Cache Writes? [yes]`
: Specify whether or not to enable writing to the cache when possible. Answering "yes" causes data to be written directly to the cache, which subsequently sends the data to the back–end. Otherwise, data is written directly to the slow back–end device in pass–through mode.

`Cache Only File System Metadata? [no]`
: Metadata captures important file system statistics such as file system inodes, size, the date stamp, and owner.

  If you are caching on a raw device (there is no file system structure in place, and consequently, there is no file system metadata), answer "no" to not cache only file system metadata. Otherwise, in this instance, no data could be cached.

  You may choose to answer "yes" when caching the application's data may not be useful, but caching only the metadata is useful. Two primary instances follow. 1) A database manager application is responsible for regulating its own data transmissions. Allowing the DG/UX system to perform the caching

function would be superfluous and probably ineffective. 2) NFS servers using the cache over a LAN will also regulate their own buffering needs. Data received over a LAN would likely flood a cache's buffers. Selection of caching only the metadata would be preferred in these instances. File system data would be read and written in pass–through mode, bypassing the front–end entirely.

Caching both the file system data and its metadata would be desired for a local file system.

If the cache is to contain a non–DG/UX file system, such as a database, do not elect to cache file system metadata only. Such a selection would prevent any data caching.

`Asynchronous Write Policy: [First write]`

This operation is meaningful only if writes are being cached. You may request direct writes from the front–end to the back–end device immediately after caching it under certain conditions:

`Never` The system will never write data directly to the back–end. The advantage is that no redundant disk activity takes place, but the disadvantage is that cache performance may degrade if the front–end gets full.

`All writes` Any time a buffer is written to the cache, an asynchronous I/O flushes the buffer to the back–end device. This setting may improve performance in caches doing more reads than writes. A disadvantage of this operation is that substantial redundant disk activity may take place.

`First write` The first time a new buffer is written to the cache, an asynchronous I/O will flush it to the back–end device. This setting helps to keep the cache clean when there are a lot of buffers that are written only once such as for large sequential writes. However, it minimizes disk activity for those blocks that are written repeatedly. In general, "first write" is the best policy.

`Read Weight: [1]`

In pre–DG/UX 5.4R3.00 releases, read weight was referred to as read retention rate. The read weight is meaningful only if you cache read operations.

You may use the read weight parameter to assign a relative priority to read operations that occur in a front–end device. A larger weight gives higher priority to the specific operation. When the cache is being searched for available space, buffers with lower weights are reused first. Thus, buffers with a high weight are retained in the cache longer.

Alternatively, if multiple cache virtual disks share a front–end device, you assign a higher priority to one cache by increasing its read weight value. Conversely, you may assign a lower priority to one or more remaining caches by decreasing their read weight values. The maximum read weight is 1000.

`Write Weight: [1]`

In pre–DG/UX 5.4R3.00 releases, write weight was referred to as write retention rate. The write weight is meaningful only if you cache write operations.

You may use the write weight parameter to assign relative priorities to write operations occurring in a front–end device. A larger weight gives higher priority to the specific operation. When the cache is being searched for available space, buffers with lower weights are reused first. Thus, buffers with high weights are retained in the cache longer.

Alternatively, if multiple cache virtual disks share a front–end device, you assign a higher priority to one cache by increasing its write weight value. Conversely, you may assign a lower priority to one or more remaining caches by decreasing their write weight values. The maximum write weight is 1000.

`Search percentage: (0-100) [10] ?`

This is the percentage of the cache's front–end for the system to search when it is looking for a clean buffer or data block with the lowest weight. If no clean buffer is found, a dirty buffer with the lowest weight is flushed and reused. On the next search for a clean buffer, the search will begin where the last search ended, so eventually the entire cache will be searched. A value of zero causes the next buffer to be used. By default, when creating a cache, the search percentage is 10 percent.

Ideally, a cached disk provides a performance improvement by satisfying disk accesses using much faster memory accesses. There are two obstacles, however, that prevent disk caching from reaching this ideal level of performance:

The cache device is not large enough to contain all the data that applications will require of it; therefore, some I/O requests will require accessing the back end device for data not currently in cache. The ratio of I/O requests satisfied by the front end cache (*cache hits*) to the total number of I/O requests is called the *cache hit rate*. You want the cache hit rate, which is expressed as a percentage, to be as high as possible.

If the buffer that your application needs to read or write is not in the cache, or those available are dirty, you must wait for the data to be flushed. This waiting period is referred to as a *stall*, which unnecessarily degrades system performance. Increasing the search percentage will give the flusher a greater amount of buffer to search through when seeking dirty buffers to flush. Locating dirty buffers quicker will cause more frequent flushes, which reduces stall times. You want stalls to occur as seldom as possible. Increasing the search percentage will bring down the occurrence of stalls.

By experimenting with the various parameters, you can find ways to maximize the cache hit rate and minimize the frequency of stalls for your cached disk. Once you have created the cached virtual disk, use `Device -> Disk -> Virtual -> Cache -> List` to review performance statistics, and use `Device -> Disk -> Virtual -> Cache -> Modify` to adjust

the operating parameters. Also you may use the **nsar**(1M) command and AV/SysScope™ performance monitor to check performance statistics.

`Flusher Type: [Cyclic] ?`

If you cache reads only, the flusher is not used. If you cache writes only, the flusher is used. If you cache both reads and writes, the flusher walks sequentially through the cache. In a cache used for writing as well as reading, the cache must at some point write, or flush, the data in its buffers to disk. If your application accesses the cache during a flush, or if your application causes a flush, it will have to wait, or stall, until the flush completes. You want to minimize stalls. A cyclic flush occurs on a regular basis.

`Name for Cache: [`*`virtual-disk-name`*`]`

Specify the name for the cache. Normally, the cache assumes the name of the virtual disk that is being cached (for example, **payroll**), and the virtual disk being cached is given a new name (next prompt). However, you may give the cache and the virtual disk on the back–end device any names you please, as long as they are unique.

`New Name for Back End: [`*`virtual-disk-name`*`.be]`

Specify the new name for the cache back–end. You may accept the default, which is its former name appended with **.be**, by pressing Enter. You are renaming the virtual disk that is being cached. Typically, the name of the virtual disk being cached travels with the cache, and you must select a new name to associate with the cache back–end device. Alternatively, you may select a new name for the back end. You may assign the cache and the back end any names you please, as long as they are unique.

`Front End Virtual Disk:`

Specify the name of the virtual disk to serve as a fast, front–end device.

`Select space by: [Disk to partition and partition size]`

You create a virtual disk to be used as a front end by either partitioning an existing virtual disk or physical disk or using an existing virtual disk. Type ? for a list of choices. Refer to the section on methods for creating a virtual disk for a recap.

If you partition the NVRAM board, specify the **nvrd**() device as the disk from which to partition. If you intend to use the entire NVRAM board as a front end, select the largest block size offered, such as 3967 blocks. This number of blocks results from subtracting 33 blocks for system partitioning overhead from a total capacity of 4000 blocks. Alternatively, if a fast disk is used as a front end, as a rule–of–thumb, size the front end to be one–tenth the size of the back end.

`Specify Another Front End? [yes]`

Specify whether or not you want to add another front–end device to the cache. If you answer "yes," the preceding prompt repeats. Otherwise, no more prompts are delivered, and the new cache is built according to your selections.

NVRAM Caching and Continuous Data Access: A Typical Scenario

Suppose you have a file system on a local AVIiON server which you later decide to export over the network. The drawback of exported file system I/O is its slow speed. One of the principal causes of slowness of remotely mounted file systems is delayed data write confirmations. To maintain data integrity, ONC/NFS does not acknowledge a successful write operation until the data is actually written to the remote physical device.

Caching is one answer to reducing write confirmation delays. With normal operating system circumstances (e.g., sufficient CPU), the ability to dynamically insert a cache encourages you to experiment with your configuration with no inconvenience due to down time. Through experimentation, you decide whether or not to cache the virtual disk.

Assume that you have a partition named **notes** that resides on a slow, back–end device for which you want to create a fast front–end device. The following table identifies the goals for this cache.

Cache name	Back–end	Size in blocks	Front–end	Size in blocks
notes	**notes–be**	100,000	**notes–fe**	3,967

A typical **sysadm** session to cache virtual disk **notes** follows.

Before you start, realize that the virtual disk to be cached is named **notes**. When you create the cache for the virtual disk, the cache will inherit the name of the back–end virtual disk, and the back–end virtual disk must take on a new name. For example, when caching the **notes** virtual disk, the cache will inherit the virtual disk's name, **notes**, and the back–end must assume a new name, such as **notes–be**. Naturally, the front–end of the **notes** cache would be **notes–fe**. The cache itself, by convention, assumes the name of the virtual disk being cached.

To create a front–end virtual disk, **notes–fe**, perform the **sysadm** `Device -> Disk -> Virtual -> Create` operation.

To build a cache from existing virtual disks, select the **sysadm** operation `Device -> Disk -> Virtual -> Caches -> Cache`.

```
Virtual Disk: notes↵
Cache Reads? [yes] no↵
Cache Writes? [yes] ↵
Cache Only File System Metadata? [no] ↵
Asynchronous Write Policy: [First write] ↵
Read Weight: [1] ↵
Write Weight: [1] ↵
Search Percentage: (0-100) [10] ↵
Flusher Type: [Cyclic] ↵
```

```
New Name for Cache: [notes] notes ↵
New Name for Back End: notes-be ↵
Specify a Front End? [no] yes ↵
Front End Virtual Disk: notes-fe ↵
Specify Another Front End? [yes] no ↵
Virtual disk inserted at notes.
Child virtual disk made a volume.
Formatting virtual disk notes-fe for use as a cache front end device.
This will destroy any data on the virtual disk.  Procede? [no] yes ↵
Formatting virtual disk notes-fe for use as a cache front end device.
Virtual disk "notes-fe" formatted for use as a cache front end.
Virtual disk "notes-fe" linked as a child to virtual disk "notes".
```

Linking One or More Front–End Devices to an Existing Cache

To link one or more front–end devices (existing virtual disks on fast devices) to an existing cache, perform the **sysadm** operation: `Device -> Disk -> Virtual -> Caches -> Link` operation.

`Existing Cache:`
: Specify the name of the existing cache to which to link one or more front–end devices.

`Front End Virtual Disk:`
: Specify the name of the virtual disk to serve as a fast, front–end device. You must select a pre–existing virtual disk. Type ? for a list of choices.

`Specify Another Front End? [no]`
: Specify whether or not you want to add another front–end device to the cache. If you answer "yes," the preceding prompt repeats. Otherwise, no more prompts are delivered, and one or more front–end device(s) is linked to the cache, according to your selections.

Unlinking One or More Front–End Devices from an Existing Cache

To unlink (remove) one or more front–end devices from an existing cache, select the **sysadm** operation `Device -> Disk -> Virtual -> Caches -> Unlink`. When the front–end device is unlinked, all data in the front–end device is flushed to the back–end device automatically.

After you select this operation, the system displays the following queries:

`Existing Cache:`
: Specify the cache from which to unlink the front–end device.

`Front End Virtual Disk:`
: Specify the name of the front–end device to be unlinked. You must select a pre–existing virtual disk. Type ? for a list of choices.

`Specify Another Front End? [yes]`

Specify whether or not you want to unlink another front–end device from the cache. If you answer "yes," the preceding prompt repeats. Otherwise, no more prompts are delivered, and one or more front–end device(s) is unlinked from the cache, according to your selections.

Following your confirmation, the desired front–end device is unlinked (removed) from the cache. The virtual disk still exists, but is not associated with the cache.

Uncaching (Dismantling) a Cache

To dismantle a cache (unlink its front–end devices), restoring the cache's name to the back–end virtual disk and remove the cache itself, select the **sysadm** operation `Device -> Disk -> Virtual -> Caches -> Uncache`. All data cached in the front end is first flushed to the back end before the system dismantles the cache. The operation deletes the cache, leaving behind its child virtual disks which continue to be used as file systems or by the application. The former back–end device resumes its original name—the cache's name.

After you select this operation, the system displays the following queries:

`Existing Cache:`

Select the cache virtual disk to disassemble.

Following your confirmation, the desired cache is disassembled. The virtual disks still exist, but are not associated with the cache.

After you have disassembled a cache, you can use its virtual disks just as you would any virtual disks.

Modifying a Cache's Attributes

To alter any of these cache attributes: cache reads, cache writes,whether or not to cache only file system metadata, asynchronous write policy, read weight, write weight, search percentage, or flusher type, select the **sysadm** operation `Device -> Disk -> Virtual -> Caches -> Modify`.

After you select this operation, the system displays the following queries:

```
Virtual disk:

Cache Reads? [yes]

Cache Writes? [yes]

Cache Only File System Metadata? [no] ?

Asynchronous Write Policy: [First write]

Read Weight: [1]

Write Weight: [1]

Search Percentage: (0-100) [10] ?

Flusher Type: [Cyclic] ?
```

Refer back to the section on caching a virtual disk for descriptions of these prompts.

Listing Cache Statistics

To list information about caches and their front– and back–end devices, select the **sysadm** operation `Device -> Disk -> Virtual -> Caches -> List`.

After you select this operation, the system displays the following queries:

`Cached Disk(s) to List: [all]`
: Select all caches or specify by name those you wish to see.

`List Front and Back Ends? [no]`
: Answering "no" lists only the caches; "yes," the caches' constituent front–end and back–end devices as well.

When you list caches only, the information displayed is a subset of the attributes you gave the cache when you created it. Besides the name, the "Volume" value specifies whether or not it has been made into a volume: "V" represents a virtual disk that is a volume, and "L" means that the virtual disk has only a long–name entry in **/dev/dsk**. "T" means that the virtual disk is temporary and will be lost on system shutdown. "F" means the virtual disk is floating (changes to virtual disk definition will be lost on system shutdown.). It also includes the size of the cache (that is, the of its virtual disk used as a back–end device), and the back–end device name.

A typical listing for caches only follows:

```
Name                          Volume Temp  Size Type
foo                                V        100 cache for zoo,2CBDD094
home                               V        200 cache for home-back
notes                              V      40000 cache for notes-back
```

The listing for the cache and its front– and back–end devices is quite lengthy, including all the attributes you specified when creating it. A typical example follows:

```
Virtual disk name:  notes
It is usable, not in use, a volume, persistent, 40000 blocks.
Its type is:  cache for notes-back.
The cache is available.
Fraction of reads coming from a front end device (read hits): n/a
Fraction of writes going to a front end device (write hits): n/a
Caching reads, caching writes.
Not caching only file system meta-data.
First write to cache triggers an asynchronous write to the back end.
Use of direct memory access by front-end devic(s) is enabled.
Read retention weight: 1
Write retention weight: 1
Percentage of front-end(s) to search for available buffer: 10%
Flusher type:  none.
Read count: 0
Write count: 0
Read miss count: 0
Write miss count: 0
Read hit count: 0
Write hit count: 0
Read purge count: 0
```

```
Write purge count: 0
Allocate purge count: 0
Back end read count: 0
Back end write count: 0
Total buffers count: 1408
Total data blocks count: 46008
Dirty buffers count: 0
Allocated buffers count: 0
Total data blocks count: 46008
State stall count: 0
Flush stall count: 0
Cache back end:
   Virtual disk name:  notes-back
   It is usable, in use, a volume, persistent, 40000 blocks.
   Its type is:  partition on sd(insc(0),0,0).
   It starts at block 44747 and is 40000 blocks long
   Bad block mapping is enabled.
Cache front end 1:
   Virtual disk name:  notes-front
   It is usable, in use, a volume, persistent, 40000 blocks.
   Its type is:  partition on sd(insc(0),0,0).
   It starts at block 84747 and is 40000 blocks long
   Bad block mapping is enabled.
Cache front end 2:
   Virtual disk name:  foolish
   It is usable, in use, a volume, persistent, 6100 blocks.
   Its type is:  partition on sd(insc(0),0,0).
   It starts at block 18347 and is 6100 blocks long
   Bad block mapping is enabled.
```

Changing the Size of a Cached Virtual Disk

To change the size of a cached virtual disk, you must perform these explicit steps:

- Disassemble the cache, leaving intact the child virtual disk that was being cached. The child virtual disk (back–end device) inherits its parent cache's name. Use the **sysadm** operation `Device -> Disk -> Virtual -> Caches -> Uncache`.

- Expand the virtual disk appropriately using the **sysadm** operation `Device -> Disk -> Virtual -> Expand` or `File System -> Local -> Expand`.

- Re–create the cache using the front–end devices that you previously unlinked and the expanded virtual disk to be used as the back–end device. Use the **sysadm** operation `Device -> Disk -> Virtual -> Caches -> Cache` to re–create the cache.

Managing Disk Arrays

Disk array subsystems have similar, but not identical, management utilities. For a HADA-I subsystem, the utility is **Gridman** (**usr/bin/gridman**); to run it, use the **sysadm** sequence `Device -> Disk -> Manage HADA-1 Array`. For a CLARiiON disk-array storage system, the utility is **GridMgr**; it runs continuously on the storage-system console. For more information on the management utility, see the 014-series disk-array storage system manual supplied with the storage system.

End of Chapter

Chapter 8
Setting Up and Managing Features of High Availability

This chapter addresses three primary hardware and software configurations that enhance high availability. They are:

- Failover disks
- Internet Protocol (IP) takeover
- Multi–path LAN I/O failover

Managing Failover Disks

Disk failover and tape sharing involves one or more disk (or tape) drives accessible by two different routes (paths). Generally, the routes run to controllers in two hosts but sometimes they run to different controllers within the same host.

Disk failover is most useful where two host computers are connected to a common disk-array storage system; it provides higher availability of the stored data. If one host or a disk controller fails, the second system can take control of the disk by performing a *trespass* and make the data available again. Tape sharing is primarily useful because it lets you use tape drives more efficiently; both systems can use one tape drive.

A failover disk (or shared tape) can be on a shared SCSI bus in a *dual-initiator* configuration. Or with a disk-array storage system, a failover disk can be on a split bus in a cabinet-sharing configuration. Figures 8–1 though 8–3 show devices in shared-bus/dual-initiator configurations. Figure 8–1 shows two hosts with a CSS/3 subsystem and Figures 8–2 and 8–3 show two hosts with a disk-array storage system. Figure 8–4 shows two hosts with a disk-array storage system in a split-bus configuration.

NOTE: Each virtual disk mounted on a host must have a unique name. Therefore, if you want to set up a system disk for failover to another host's system disk, the virtual disks on that system disk must have different names from the virtual disks on the other host's system disk (for example, **root1**, **usr1**, and **opt1**).

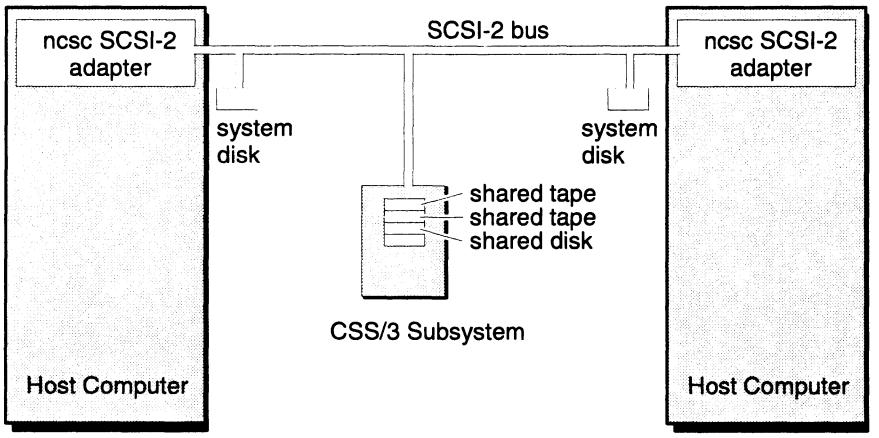

Figure 8–1 Two hosts with CSS/3 Subsystem in Shared-Bus/Dual-Initiator Configuration

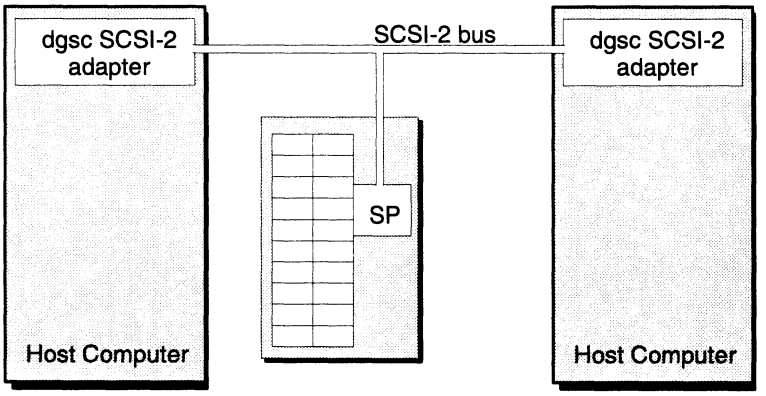

Figure 8–2 Two hosts with Disk-Array Storage System in Single Shared-Bus/ Dual-Initiator Configuration

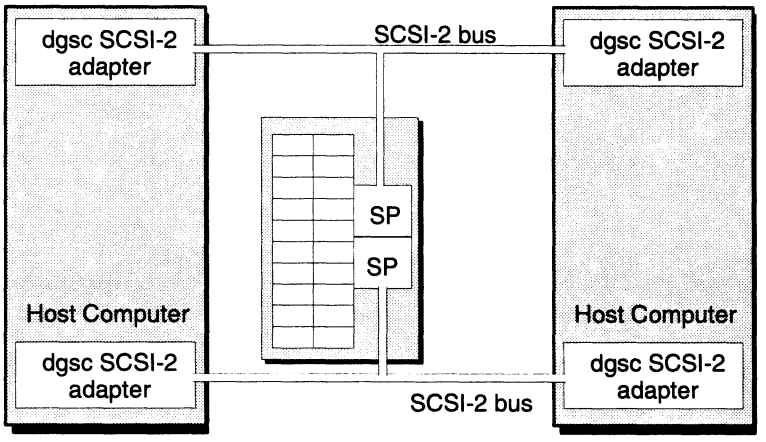

Figure 8–3 Two hosts with Disk-Array Storage System in Dual Shared-Bus/ Dual-Initiator Configuration

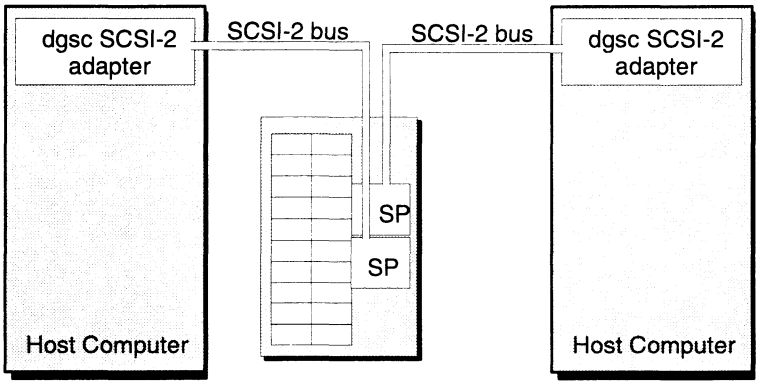

Figure 8–4 Two hosts with a Disk-array Storage System in a Split-Bus Configuration

The way you set up a device to be shared depends on the configuration you want to support. Skip to the pertinent section following.

- Setting up Devices on a Shared SCSI Bus without a Disk-Array Storage System
- Setting up Disks on a Shared SCSI Bus with a Disk-Array Storage System
- Setting up Disks on a Split SCSI Bus with Disk-Array Storage System

Shared SCSI Bus without a Disk-Array Storage System

This section tells how to set up two computers to share a SCSI bus in a shared-bus/dual-initiator configuration without a disk-array storage system; that is, with a peripherals unit like a CSS/3 subsystem. This configuration works only on Data General AViiON systems that support SCSI-2 adapters (such as **ncsc**). For both hosts to use a shared bus, each must specify a different SCSI address for its SCSI-2 adapter channel. For example, one host uses disk unit 2 though the disk unit name **sd(ncsc(0,7),2,0)** and the other host uses disk unit 2 through the disk unit name **sd(ncsc(0,6),2,0)**. The 7 and 6 in the unit names are the SCSI adapter SCSI IDs; since they are different, they allow each host to initiate its own requests to the bus.

The basic requirement for a shared SCSI bus configuration is that every device on the bus must have a different SCSI ID. This means the SCSI-2 adapters themselves must have different SCSI IDs.

CAUTION: If two or more devices on the SCSI bus have the same SCSI ID, either or both systems may hang or behave erratically. Be sure that each device on the bus, including each SCSI adapter, has a unique SCSI ID.

To set up two systems to share a SCSI bus without a disk-array storage system, follow these steps:

1. For each device that will be on the shared SCSI bus (other than the SCSI adapters themselves), set the jumpers to indicate a unique SCSI ID. Do not set any devices to SCSI ID 6 or 7; reserve these numbers for the SCSI adapters later.

For example, you want to set up hosts **system1** and **system2** to share a SCSI bus. The SCSI bus will support a combined storage subsystem with three disk drives (including a system disk for each host) and two tape drives, as shown earlier in Figure 8–1. Each device on the bus will have a unique SCSI ID, as follows.

Device	SCSI ID
1.4 Gbyte disk (system disk)	0
1.4 Gbyte disk (system disk)	1
1.4 Gbyte disk (to be shared)	2
QIC-320 tape (to be shared)	4
QIC-320 tape (to be shared)	5

2. Install the hardware: connect the devices to either of the two systems in the standard fashion, such that you can boot and set up both systems independently. The disk that you intend to use for failover may be connected to either system. Do not connect the two systems' SCSI buses together at this time.

   In the example configuration, you connect the systems so that **system1** has two disks and one tape drive, and **system2** has one disk drive and one tape drive, as follows.

system1	system2
`sd(ncsc(),0,0)`	`sd(ncsc(),1,0)`
`sd(ncsc(),2,0)`	`st(ncsc(),5,0)`
`st(ncsc(),4,0)`	

   The second disk attached to **system1**, the one at SCSI ID 2, will be the failover disk.

3. Select one of the systems to set up first. Power up and boot the system. If necessary, install the DG/UX system software. This is the system whose SCSI-2 adapter will have the SCSI ID of 6. You set the adapter's SCSI ID when you build the kernel, discussed below.

4. Execute the **sysadm** sequence `System -> Kernel -> Build` to build a new kernel. For any tape drive you want the hosts to share, make sure the device entry appears in both host's system files. Do not, however, add an entry for a failover disk to both systems.

CAUTION: Include the failover disk device entry only in the system file of the system where it is installed. If you include the device entry for a failover disk in the system files for both systems, the systems will enter a race condition upon booting, and only one of the systems, the one that configures and registers the disk first, will be able to access the disk.

Make sure the entries in the system file include the correct SCSI ID for the device as well as the SCSI ID for the local SCSI adapter, SCSI ID 6. Build and install the kernel.

In the example configuration, boot **system1** and execute `System -> Kernel -> Build`. When **vi** runs, you see the following entries in the system file:

```
sd(ncsc(),0,0)
sd(ncsc(),2,0)
st(ncsc(),4,0)
```

You want let the hosts share the tape drive attached to **system2**, so you add an entry for it to **system1**'s system file:

```
st(ncsc(),5,0)
```

Then add the SCSI adapter address and SCSI ID in all of these entries. The SCSI-2 adapter number is 0 (it is the first such adapter in the host), and its SCSI ID is 6. The list in the system file now looks like this:

```
sd(ncsc(0,6),0,0)
sd(ncsc(0,6),2,0)
st(ncsc(0,6),4,0)
st(ncsc(0,6),5,0)
```

5. After making the changes you want to the kernel configuration, exit from the text editor (for **vi**, enter Esc, then ZZ).
6. Proceed with the Build operation and install the new kernel. If an error occurs, **sysadm** will display an error message. For more information on building a kernel, see Chapter 4.
7. Shut down the first system. From the SCM, change the SCM's default boot path so that it includes the correct SCSI adapter SCSI ID. For example, if the default boot path had been **sd(ncsc(),0,0)**, change it to **sd(ncsc(0,6),0,0)**. If you also intend to set your system's boot path using **dg_sysctl**(1M) after you have brought the system back up, remember to specify the correct SCSI adapter SCSI ID there as well.

   CAUTION: Be careful to specify the correct SCSI adapter SCSI ID in the boot path. If the boot path you use to boot your system specifies the SCSI ID of the SCSI adapter installed in the other system, the boot will fail and the SCSI bus will hang. If the SCSI bus hangs, an attempt by either system to access the SCSI bus will hang the system. When this happens, recover by resetting the hardware and rebooting.

8. After resetting the SCM boot path, power off the system.
9. Power on and boot the second system. If necessary, install the DG/UX system software. This system's SCSI adapter will have SCSI ID 7.
10. Execute `System -> Kernel -> Build` and perform essentially the same steps that you performed for the first system: add device entries for any tape drives connected to the other system that you want to share; then in each device entry set the SCSI ID for the SCSI adapter. On this system, the SCSI adapter SCSI ID is 7.

For example, you execute `System -> Kernel -> Build`. The system file initially contains these entries:

```
sd(ncsc(),1)
st(ncsc(),5)
```

You want to share the tape drive connected to **system1**, so you add the entry for it in your system file:

```
st(ncsc(),4)
```

Then add the SCSI adapter address and SCSI ID in all of these entries. The SCSI adapter address is 0, and the SCSI ID is 7. The list in the system file now looks like this:

```
sd(ncsc(0,7),1,0)
st(ncsc(0,7),5,0)
st(ncsc(0,7),4,0)
```

11. After making the changes you want to the kernel configuration, exit from the text editor (for **vi**, enter Esc, then ZZ).

12. Proceed with the Build operation and install the new kernel. If an error occurs, **sysadm** will display an error message. For more information on building a kernel, see Chapter 4.

13. Shut down the second system. Once more, from the SCM, change the SCM's default boot path so that it includes the correct SCSI adapter SCSI ID. For example, if the default boot path had been **sd(ncsc(),1,0)**, change it to **sd(ncsc(0,7),1,0)**. If you also intend to set your system's boot path using **dg_sysctl**(1M) after you have brought the system back up, remember to specify the correct SCSI adapter SCSI ID there as well.

    CAUTION: Be careful to specify the correct SCSI adapter SCSI ID in the boot path. If the boot path you use to boot your system specifies the SCSI ID of the SCSI adapter installed in the other system, the boot will fail and the SCSI bus will hang. If the SCSI bus hangs, an attempt by either system to access the SCSI bus will hang the system. When this happens, recover by resetting the hardware and rebooting.

14. After resetting the SCM boot path, power off the system.

15. With both machines powered off, reconnect the hardware in the desired configuration: using normal SCSI cables, connect the two systems and the devices as a single SCSI bus. Unlike a typical SCSI bus, this bus has no external terminators on it; the adapters in the systems at each end terminate the bus. Refer to your hardware documentation before performing this step.

16. Power on and reboot the first system with its new kernel.

17. When the first system has completed booting, power on and reboot the second system with its new kernel.

Table 8–1 shows the initial configuration of the example systems.

Table 8–1 Sample Shared-Bus Configuration without Disk-Array Storage System

Device	SCSI ID Number	Name Relative to system 1	Name Relative to system 2
SCSI adapter on **system 1**	6	**ncsc(0,6)**	Not accessible
SCSI adapter on **system 2**	7	Not accessible	**ncsc(0,7)**
1.4 Gbyte disk (system disk)	0	**sd(ncsc(0,6),0,0)**	Not accessible
1.4 Gbyte disk (system disk)	1	Not accessible	**sd(ncsc(0,7),1,0)**
1.4 Gbyte disk (shared)	2	**sd(ncsc(0,6),2,0)**	Normally not accessible
QIC 320 tape drive (shared)	4	**st(ncsc(0,6),4,0)**	**st(ncsc(0,7),4,0)**
QIC 320 tape drive (shared)	5	**st(ncsc(0,6),5,0)**	**st(ncsc(0,7),5,0)**

At this point, users on the two systems may share both tape drives. Initially, the only **system1** can access the failover disk (SCSI ID 2). To enable failover for this disk, you must use **sysadm** to configure the disk for operator-initiated failover (OIF). You will do this by adding the disk to the **sysadm** giveaway database (`Availability -> Disk Failover -> Giveaway -> Add`) on both hosts, as described later in this chapter, section "Using Failover Disks." You can then use **sysadm** to transfer the disk from **system1** to **system2**.

Operating a Shared Tape Drive

Sharing a tape drive is essentially the same in a dual-initiator configuration as on a single system: the tape drive is available for anyone until someone opens it, at which time it remains occupied until the user closes it. The commands that you typically use to open a tape drive include **cpio**, **tar**, and **dump2**. Any attempt to open an already-opened shared tape drive results in the same error that you might see when two users on the same system try to use an unshared tape drive at the same time. Thus users take turns in sharing a tape drive on the shared bus the same way they take turns when sharing a normal one.

If either system boots while the other is using a tape device on the shared SCSI bus, I/O operations to the tape may fail. If a user or application on the running system is using a shared tape device when the other system boots, one of two things will occur: the user or application using the tape drive will receive an I/O error; or the booting system will be unable to configure the tape device. To avoid these problems, make sure no shared tape drives is in use when booting either system.

Shared SCSI Bus with a Disk Array Storage System

This section tells how to set up two computers to share a SCSI bus in a shared-bus/dual-initiator configuration. It applies only if you have a disk-array storage system with an adapter that supports SCSI–2 protocol (such as a **dgsc**). For both hosts to use a shared bus, each must specify a different SCSI ID for its SCSI-2 adapter channel. For example, one host uses disk unit 0 via the disk unit name **sd(dgsc(0,7),0,0)** and the other host uses disk unit 1 via the disk unit name **sd(dgsc(0,6),0,1)**. The 7 and 6 in the unit names are the SCSI-2 adapter SCSI IDs; they allow each host to initiate its own requests to the bus.

The basic requirement for a shared SCSI bus configuration is that every device on the bus must have a different SCSI ID. This means the SCSI-2 adapters themselves (one or two on each system), must have different SCSI IDs as well. Essentially, with a disk-array storage system, the disk units themselves do not have SCSI IDs; instead, the SPs (system control processors) have SCSI IDs. Therefore in the disk unit names shown above, if you omit the LUN numbers at the end, the SP device names become the unique names **sd(dgsc(0,7),0)** and **sd(dgsc(0,6),0)**; this satisfies the requirement for a unique SCSI ID.

CAUTION: If two or more devices on the SCSI bus have the same SCSI ID, either or both systems may hang or behave erratically. Be sure that each device on the bus, including each SCSI adapter, has a unique SCSI ID.

To set up two systems to share a SCSI bus, follow these steps:

1. If there are two SPs, make sure each is set to a unique SCSI ID (usually 0 and 1). The default SCSI ID for the first SP (SP A) is 0; for the second SP (SP B), it is 1. Do not set any devices to SCSI ID 6 or 7; reserve these numbers for the the SCSI-2 adapters.

   For example, you are setting up systems **host1** and **host2** to share a SCSI bus. The SCSI bus will connect to an SP that runs three disk drives on one SP, SP A, SCSI ID 0. The device name of the SP from each host will need to be unique. The disks, the SCSI-2 adapters, SPs, and the disk drive names on the two hosts are as follows in Table 8–2 .

Table 8–2 Sample Shared-Bus Dual-Initiator Configuration with Disk-Array Storage System

Device	SCSI ID		Physical unit number (LUN)	Disk Drive Name Relative to host1	Disk Drive Name Relative to host2
	Adapter	SP			
SCSI-2 adapter on **host1**, channel 0	6	0	Does not apply	**dgsc(0,6)**	Not accessible
SCSI-2 adapter on **host2**, channel 0	7	0	Does not apply	Not accessible	**dgsc(0,7)**
1.0 Gbyte disk (**host1** system disk)	6	0	0	**sd(dgsc(0,6),0,0)** (name for **host1** system file)*	Not accessible
1.0 Gbyte disk (**host2** system disk)	7	0	1	Not accessible	**sd(dgsc(0,7),0,1)** (name for **host2** system file)*
4.0 Gbyte disk (failover disk)	7	0	2	**sd(dgsc(0,6),0,2)** (name for **host1** system file)*	**sd(dgsc(0,7),0,2)** (path for failover, not in system file)

* For a host that has a second VME channel, the device name must also specify the VME channel. The VME channel name has the form **vme(*n*)**, and immediately follows the SCSI-2 adapter name. For example, if **host1**'s system disk is on the second VME channel, **vme(1)**, the **host1** system disk device name is **sd(dgsc(vme(1),0,6),0,0)**. The **probedev** autosizer program automatically generates **vme** entries in the system file for you.

2. Install the hardware: connect the devices to either of the two systems in the standard fashion, such that you can boot and set up both systems independently.

3. Select one of the systems to set up first. Power up and boot the system. If necessary, install the DG/UX system software. This is the system whose SCSI-2 adapter will have the SCSI ID of 6.

4. Execute the **sysadm** operation `System -> Kernel -> Build` to build a new kernel. Proceed as usual. When **vi** runs, and you look for disk controllers in the system file, you will see lines that look something like this:

```
sd(dgsc(0),0)    ## SCSI disk 0 on Data General SCSI adapter 0
sd(dgsc(1),1)    ## SCSI disk 1 on Data General SCSI adapter 1
```

5. Make sure the entries in the system file include the correct SCSI-2 adapter and SCSI ID, correct SP SCSI ID, and correct disk drive unit number for each disk. For this host, this adapter, and SP, you want the system file to include the the **host1** system disk and failover disk:

   sd(dgsc(0,6),0,0)
   sd(dgsc(0,6),0,2)

   So for **host1** you would make sure the sd lines for that disk-array storage system look like this, adding the comments for clarity:

```
sd(dgsc(0,6),0,0)    ## 1st SCSI-2 adapter, 1st channel (0); SP A (ID 0); unit 0
                     ## Host1 system disk.
sd(dgsc(0,6),0,2)    ## 1st SCSI-2 adapter, 1st channel (0); SP B (ID 0); unit 2,
                     ## failover disk.
```

- For a shared or split bus configuration, never use the asterisk notation for SCSI devices [**sd(dgsc(0),***]. Specify the entire, specific device names as explained above. If two hosts sharing a SCSI bus use the asterisk to configure all devices on the bus, the first host to start up will assume control of all devices, and the other will not be able to use any.

6. After making the changes you want to the kernel configuration, exit from the text editor (for **vi**, enter Esc, then ZZ).

7. Proceed with the Build operation and install the new kernel. If an error occurs, **sysadm** will display an error message. For more information on building a kernel, see Chapter 4.

8. Shut down the first system. While it is in the SCM, change the SCM's default boot path so that it includes the correct SCSI adapter SCSI ID. For example, if the default boot path had been **sd(dgsc(),0,0)**, change it to **sd(dgsc(0,6),0,0)**. If you also intend to set your system's boot path using **dg_sysctl**(1M) after you have brought the system back up, remember to specify the correct SCSI adapter SCSI ID there as well.

   CAUTION: Be careful to specify the correct SCSI adapter SCSI ID in the boot path. If the boot path you use to boot your system specifies the SCSI ID of the SCSI adapter installed in the other system, the boot will fail and the SCSI bus will hang. If the SCSI bus hangs, an attempt by either system to access the SCSI bus will hang the system. When this happens, recover by resetting the hardware and rebooting.

9. After resetting the SCM boot path, power off the system.

10. Power on and boot the second system. If necessary, install the DG/UX system software. This system's SCSI adapter will have SCSI ID 7.

11. Execute `System -> Kernel -> Build` and perform essentially the same steps that you performed for the first system: add device entries for any tape drives connected to the other system that you want to share; then in each device entry set the SCSI ID for the SCSI adapter. On this system, the SCSI adapter SCSI ID is 7.

    For example, in **vi**, the sample entries look like this:

```
sd(dgsc(0),0)    ## SCSI disk 0 on Data General SCSI adapter 0
sd(dgsc(1),1)    ## SCSI disk 1 on Data General SCSI adapter 1
```

    And for **host2**, you will change these to

```
sd(dgsc(0,7),0,1)    ## 1st SCSI-2 adapter, 1st channel (0); SP A (ID 0); unit 0
                     ## Host2 system disk.
```

    And build and install the kernel as on the other host.

12. Shut down the second system. From the SCM, change the SCM's default boot path so that it includes the correct SCSI adapter SCSI ID. For example, if the default boot path is **sd(dgsc(),1,0)**, change it to **sd(dgsc(0,7),1,1)**. If you also intend to set your system's boot path using **dg_sysctl**(1M) after you have brought the system back up, remember to specify the correct SCSI adapter SCSI ID there as well.

    CAUTION: Be careful to specify the correct SCSI adapter SCSI ID in the boot path. If the boot path you use to boot your system specifies the SCSI ID of the SCSI adapter installed in the other system, the boot will fail and the SCSI bus will hang. If the SCSI bus hangs, an attempt by either system to access the SCSI bus will hang the system. When this happens, recover by resetting the hardware and rebooting.

13. After resetting the SCM boot path, power off the system.

14. Power on and reboot the first system with its new kernel.

15. When the first system has completed booting, power on and reboot the second system with its new kernel.

You have set up the configuration shown in Table 8–2.

At this point, only **host1** can access the failover disk on SCSI ID 2. To enable **host2** to take over this disk, you must use **sysadm** to configure the disk for operator-initiated failover (OIF) and then optionally for machine-initiated failover (MIF). To enable OIF, you will do this by adding the disk to the **sysadm** giveaway database on both hosts, as described later in this chapter, section "Using Failover Disks." You can then use **sysadm** to transfer the disk from **host1** to **host2** and, if you want, enable machine-initiated failover.

Split SCSI Bus with a Disk Array Storage System

This section tells how to set up two computers to share a SCSI bus in a split bus configuration, which when used with two hosts is also known as the cabinet-sharing configuration. You can use this only if you have a disk-array storage system with an adapter that supports SCSI-2 protocol (such as a **dgsc**).

A split-bus/cabinet-sharing configuration actually has two busses, each connected to its own SP. A split bus has one important advantage over a shared (dual-initiator) bus. Since the split bus uses a path of entirely different components, it offers higher availability than a shared bus. For example, if the SP in a shared-bus path fails, all access to the SP's disks is lost; if an SP in a split-bus path fails, the system can route I/O though the other SP.

The single-host and dual-host split-bus configurations are shown in Figure 8–5.

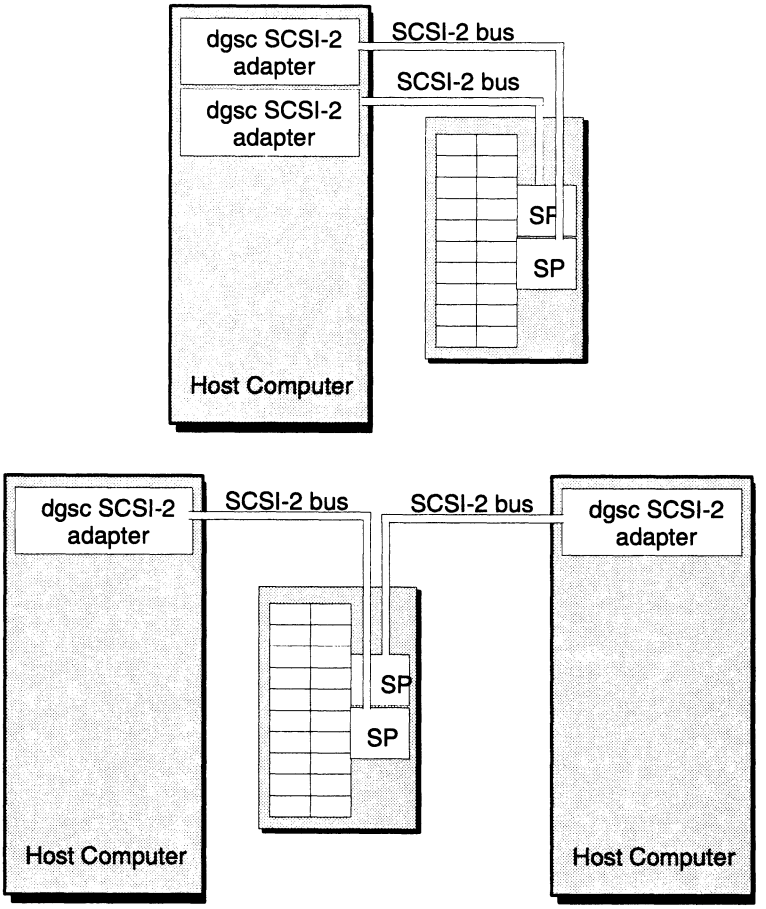

Figure 8–5 Split-Bus Configuration in One Host and Two Hosts

When you generate a kernel (in one host or two) to support a split-bus configuration, you must fully qualify the disk unit names in the system file so that each host (or path) registers only the disks you want it to. For example, one host uses disk unit 0 via the disk unit name **sd(dgsc(0),0,0)** and the other host uses disk unit 1 via the disk unit name **sd(dgsc(0),0,1)**. (You do not need to specify the SCSI-2 adapter SCSI IDs as with a shared bus.)

For more information on setting up and generating kernels to support a split-bus configuration, see the the 014-series manual that accompanies the disk-array storage system hardware.

Using Failover Disks

The section explains what failover does and the tasks involved in managing failover disks. You can implement disk failover with any disk on a shared SCSI-2 bus or with any disk on a split bus. Disk failover is most useful with disks in a disk-array storage system.

The DG/UX failover feature allows you to connect a physical disk to two systems and transfer control of the disk from one system to the other as needs arise. The failover feature not only handles operations involved in shutting off access to the disk on one system, but it also handles operations involved in making the disk accessible on the other system and starting any applications used to access it.

The systems that will be using the failover disks must be connected by a TCP/IP network. The failover daemons on the systems communicate over a port running a SAF **listen** port service added during setup of the DG/UX system software.

Failover Types

There are two types of disk failover: operator-initiated failover (OIF), which requires human intervention, and machine-initiated failover (MIF) which occurs without human intervention. To set up OIF, you need two hosts and a disk-array storage system (with one or two SPs) or one host and two SPs.

To set up MIF, you need OIF set up, and you need two hosts, each ideally with two TCP/IP network interfaces connected to the other host. With MIF, a failover monitor daemon process runs in a host and periodically checks the other host to make sure it is still running. If the other host is not running, the monitoring host takes over the disks specified in the giveaway database (set up via **sysadm**). Then the monitoring host runs a script whose name you specified; this script can restart applications and issue whatever other shell commands are needed to restore the production environment. When the surviving host detects that the other host is running again, it runs a regain-pulse script whose name you specified; this script can unmount and deregister the disks it took over so that the original owning host can mount and use them.

MIF will work with only one TCP/IP network interface, but if the network fails, the monitoring host will assume that the other has failed and will take over the other host's disks.

What Failover Does

Only one system may configure and register a physical disk at a time; therefore, a failover disk can be accessible only on one system at a time. When you use an OIF procedure to transfer control of a failover disk from one system to another, the system performs several tasks.

If the system that normally uses the failover disk fails, a human operator or the MIF software performs the failover from the other host by *trespassing*. Trespassing forces disk ownership to change. The system taking over the disk

1. Issues the command needed to transfer control of the disk.
2. Configures the disk.
3. Registers the disk.
4. Runs **fsck**(1M) to check file systems.
5. Mounts and exports file systems.
6. Starts applications used to access the disk.
7. Logs and reports operation status.

If the system giving the failover disk is running (as will often not be true), it

1. Terminates processes accessing the disk.
2. Unmounts file systems on the disk.
3. Deregisters the disk.

The effect of these tasks varies depending on how you have configured the systems. For example, if you have not built file systems on the disk, the steps related to file systems do not occur.

Depending on how you have configured the disks on your system, failover may involve multiple physical disks. This is true if the virtual disks or file systems designated for failover span multiple physical disks, as in the following cases:

- A file system was built on a virtual disk whose components span multiple physical disks.
- A striped virtual disk, which necessarily resides on multiple physical disks.
- A software-mirrored virtual disk whose images reside on multiple physical disks.

In any of the cases above, the system requires that all such interdependent physical disks be failed over together. The system does not allow you to perform failover such that only part of a virtual disk or mirror is being moved to the other system; either all parts of the virtual disk or mirror failover, or the operation fails.

Depending on how you have set up your system, there may also be other cases where a failover involves multiple physical disks. For example, you may have a database application in a file system on one physical disk and the database itself located on another physical disk. In this case, you set up both disks for failover and move them together.

The **sysadm** operations that support the failover feature are beneath the `Availability -> Disk Failover` menu. The **sysadm** operations call the **admfailoverapplication**(1M), **admfailoverdisk**(1M), **admfailovergiveaway**(1M), **admfailoverhosts**(1M), and **admfailovertakeaway**(1M) commands. The manual pages provide detailed information.

OIF Failover Databases Maintained by the DG/UX System

DG/UX uses the following databases. You create and maintain them using the **sysadm** menu sequences `Availability -> Disk Failover -> Giveaway` or equivalent **sysadm** failover commands.

- **/etc/failover/application**

  The **application** database contains information about application start-up scripts to be executed when a physical disk is failed over. The full pathname of the script must be included. Any number of scripts may be included in **application**.

- **/etc/failover/hosts**

  The **hosts** database contains information about hosts in dual-initiator configurations with this host. Entries contain the name of the host, the communication path to the host ("network" is the only path currently supported), and the current synchronization status of the host. The host must be INSYNC before it can participate in a **give** or **take** operation.

- **/etc/failover/giveaway**

  The **giveaway** database contains information about the physical disks that can be given to other hosts in the dual-initiator configuration. Entries contain any and all file system information that was gathered at the time the disk was added. This includes the name of the host that is giving away the disks, information about the virtual disks on the local and remote systems, and information about the file system to unexport, unmount, and delete during a Give operation.

  Some operations that you use to maintain the **giveaway** database, Add and Delete, offer the option of synchronizing databases. If you do not synchronize databases, the system flags them as out of sync, or NOTSYNC, and will not allow you to transfer control of any failover disk. As an alternative to the Synchronize Database option in the Add, Modify, and Delete operations, you can synchronize databases with the Synchronize operation. A later section discusses the Synchronize operation.

- **/etc/failover/takeaway**

  The **takeaway** database contains information about the physical disks that can be taken from other hosts in the dual-initiator configuration. The entries are similar to those in the **giveaway** database.

For more information about these databases, see the **failover**(4M) manual page.

Setting Up and Using OIF

To set up Operator Initiated Failover (OIF) for any disk, you must execute the following steps.

1. On the host that will be the primary owner of the disk, use **sysadm** to make sure the disk is configured and registered.

2. Still on the host that will be the primary owner of the disk, with the network running, add the disk as a failover disk with the **sysadm** sequence `Availability -> Disk Failover -> Giveaway -> Add` operation. When you have added all the disks you want to use for failover, be sure to answer **yes** to the synchronize prompt. The synchronize step enters the disk name in the takeaway database of the other host, allowing the disk to be failed over. A later section discusses the Add operation in more detail.
3. With the failover disk(s) added and databases synchronized, you can transfer control of the disk(s) should the need arise. There are several ways to perform failover, as follows.
   - When the giving system has failed, you can use the taking system to transfer control to the taking system with the **sysadm** sequence `Availability -> Disk Failover -> Take` and the **Use Trespass** option.
   - When the giving system is functioning normally, you can transfer control using either system:

     from the giving system using the **sysadm** sequence `Availability -> Disk Failover -> Give`; **or**

     from the taking system using the sequence `Availability -> Disk Failover -> Take` *without* the **Use Trespass** option.

After setting up OIF as explained in this section, you can then think about setting up MIF, described in a later section.

Managing the giveaway Database

The **sysadm** sequence `Availability -> Disk Failover -> Giveaway` enables you to add, delete, modify, and list entries in the **giveaway** failover database. The following sections describe each of these operations.

Adding a giveaway Entry

Before you can transfer control of a disk from your system to another, you must add the disk to the **giveaway** database using the `Availability -> Disk Failover -> Giveaway -> Add` operation. This operation examines the contents of the specified physical disk looking for virtual disks, disk mirrors, and export information. This information is then formatted and added to the **giveaway** database. When you synchronize, the information is added to the other host's **takeaway** database.

The Add operation queries you for the following information:

`Local Physical Disk`
: Enter the name of a local physical disk, for example, **sd(dgsc(0,6),1,0)**, that is currently configured and registered.

`Remote Physical Disk`
: Enter the name that the physical disk will have on the remote system; for example, **sd(dgsc(0,7),1,0)**.

`Hostname to Fail Disk over to`
: Enter the host name of the remote system. The system must be currently accessible on the network. If necessary, **sysadm** adds this name to the **/etc/failover/hosts** database.

`Application Command Line`
: Enter the command name to be executed on the remote system after transferring the failover disk. Typically, this command line starts the application used on the remote host to access the data on the failover disk. If necessary, **sysadm** adds this command line to the **application** database.

`Synchronize database`
: Select this option to copy the information for this failover disk to the **takeaway** database on the remote system. If the remote system is inaccessible, synchronization will fail. When the remote system becomes accessible again, synchronize databases with the Synchronize operation, discussed later. Until you synchronize databases, the system considers the local giveaway database as out of sync, or NOTSYNC, and restricts you from giving *any* failover disks to other systems.

The **giveaway** database includes information normally associated with a local file system, for example, export options, **fsck** pass number, and so on. The Add operation derives this information from the local file system table (**fstab**(4)), which you may review with the `File System -> Local Filesys -> List` operation. To change this information in the giveaway database, select `Availability -> Disk Failover -> Giveaway -> Modify`, discussed later.

The Add operation scans the failover disk for dependencies on other physical disks. A failover disk is dependent if it contains any virtual disk with a part residing on another physical disk. If any such dependencies exist, the Add operation displays a warning. You cannot transfer control of a failover disk without also transferring any codependent disks with it.

Adding the disk(s) you want to use for failover is the essential software task required for OIF. After adding the disk(s) you want, you can then think about setting up MIF, described in a later section. To maintain the giveaway and other OIF databases, continue in this section.

Deleting a giveaway Entry

To remove a failover disk from the local **giveaway** database, select the operation `Availability -> Disk Failover -> Giveaway -> Delete`. If this physical disk has dependencies on other physical disks, the operation will warn you. A dependent physical disk is one that contains part of a virtual disk that also has parts residing on other physical disks. You cannot failover a disk without also failing over any disks that have dependencies on it.

The Delete operation queries you for the following information:

`Host name`
: Select the name of the remote host designated for the failover disk.

`Disk name`
: Select the name of the local physical disk designated as the failover disk.

`Virtual Disk name`
: Select the name of the virtual disk you want to delete.

`Synchronize database`
: This option removes the information for this failover disk from the takeover database on the remote system. If the remote system is inaccessible, synchronization will fail. When the remote system becomes accessible again, synchronize databases with the Synchronize operation, discussed later. Until you synchronize databases, the system considers the local giveaway database as out of sync, or NOTSYNC, and restricts you from giving *any* failover disks to other systems.

Modifying a giveaway Entry

To change the information about a failover disk in your **giveaway** database, select the operation `Availability -> Disk Failover -> Giveaway -> Modify`. You may, for example, wish to change some of the file system mounting information (such as mount point, export options, and so on) associated with a file system on the failover disk. The Modify operation queries you for the following information:

`Host name`
: Select the remote host designated for the failover disk.

`Disk name`
: Select the physical disk used for failover.

`Virtual Disk name`
: Select the pathname of the virtual disk's special file. For example, the special file for a virtual disk named **db1** would be **/dev/dsk/db1**.

`Synchronize database`
: This option updates the information about this failover disk in the takeover database on the remote system. If the remote system is inaccessible, synchronization will fail. When the remote system becomes accessible again, synchronize databases with the Synchronize operation, discussed later. Until you synchronize databases, the system considers the local giveaway database as out of sync, or NOTSYNC, and restricts you from giving *any* failover disks to other systems.

Sysadm then allows you to change any of the other information associated with the virtual disk entry, including the remote disk name and the file system mounting information. For complete discussion of file system mounting information, see the discussion of local file systems in Chapter 9.

Finally, **sysadm** offers the Synchronize database option, which you select to copy changes you made to the takeover database on the remote system. If the remote system is inaccessible, synchronization will fail. When the remote system becomes accessible again, synchronize databases with the Synchronize operation, discussed later. Until you synchronize databases, the system considers the local **giveaway** database as out of sync, or NOTSYNC, and restricts you from giving *any* failover disks to other systems.

Listing giveaway Entries

To display information about local failover disks, select the `Availability -> Disk Failover -> Giveaway -> List` operation. The display includes information for failover disks on the local system, stored in the **giveaway** database.

The entries in the **giveaway** database include the following information for each virtual disk piece on each failover disk:

`Hostname`
: The remote host designated for failover.

`Local Diskname`
: The local name of the failover disk on which the piece resides.

`Remote Diskname`
: The remote name of the failover disk on which the piece resides.

`File System Source`
: The pathname of the special file for the virtual disk, for example, **/dev/dsk/db1**.

`<FS INFO>`
: The information used for mounting the file system, such as mount point directory, **fsck** pass number, and so on. For a discussion of this information, see the discussion of local file systems in Chapter 9.

Managing the takeaway Database

The **sysadm** sequence `Availability -> Disk Failover -> Takeaway` enables you to delete, modify, and list entries in the **takeaway** failover database as follows. The other host does not need to synchronize after you change the takeaway database.

Deleting a takeaway Entry

To remove a failover disk from the local **takeaway** database, select the operation `Availability -> Disk Failover -> Takeaway -> Delete`.

The Delete operation queries you for the following information:

Host name
: Select the name of the remote host you want to delete.

`Disk name`
: Select the name of the physical disk module or modules for which you want to delete **takeaway** entries.

`Virtual Disk name`
: Select the name of the virtual disk entry you want to delete.

Modifying a takeaway Entry

To change the information about a failover disk in the **takeaway** database, select the operation `Availability -> Disk Failover -> Takeaway -> Modify`. The entries in this file are usually taken from the remote system's **giveaway** database and do not need to be modified. However, there are times when you might need to modify **takeaway** entries.

For example, if you want to failover a system disk, you would need to modify the **takeaway** database. On your system the **giveaway** entries would have file system information for / and **/usr** file systems. These get copied to the **takeaway** database on the remote host. However, a / and **/usr** already exist on this system. You could then modify the **takeaway** entry on the remote system to give the file systems different names, such as **/root_failover** and **/usr_failover**. Now the file systems would get mounted as normal file systems which enables you, for example, to correct a corrupted system file.

The Modify operation queries you for the following information:

`Host name`
: Select the host you want to modify.

`Disk name`
: Select the physical disk module or modules for which you want to modify **takeaway** entries.

`Virtual Disk name`
: Enter or select the name of the virtual disk you want to modify.

The operation then allows you to change any of the other information associated with the virtual disk entry, including the remote disk name and file system mounting information. For detailed file system mounting information, see the discussion of local file systems in Chapter 9.

Listing takeaway Entries

To display information about remote failover disks, select the `Availability -> Disk Failover -> Takeaway -> List` operation. The display includes information for failover disks on the remote system, stored in the **takeaway** database. The entries in the **takeaway** database include the following information for each virtual disk piece on each failover disk:

`Hostname`
: The remote host designated for failover.

`Local Diskname`
: The local name of the failover disk on which the piece resides.

`Remote Diskname`
: The remote name of the failover disk on which the piece resides.

`File System Source`
: The pathname of the special file for the virtual disk, for example, **/dev/dsk/db1**.

`<FS INFO>`
: The information used for mounting the file system, such as mount point directory, **fsck** pass number, and so on. For a discussion of this information, see the discussion of local file systems in Chapter 9.

Managing the application Database

The **sysadm** sequence `Availability -> Disk Failover -> Application` menu enables you to add, delete, modify, and list entries in the **application** failover database, as follows.

Adding an application Entry

Use this to enter application scripts to be executed when a disk is failed over. You can use these scripts to minimize a system's down time by restarting applications and similar tasks. To add an entry to the **application** database, select the `Availability -> Disk Failover -> Application -> Add` operation. The Add operation queries you for the following information:

`Host Name`
: Select the host containing the physical disk for which you want to add an application script.

`Disk Name`
: Select the physical disk module or modules for which you want to add an application script.

`Remote Physical Diskname`
: Enter the name of physical disk on the remote system.

`Application Command Line`
: Enter the full pathname of a script or command to be executed when the local physical disk is failed over.

Deleting an application Entry

To remove an entry in the **application** database, select the `Availability -> Disk Failover -> Application -> Delete` operation. You can either delete specific entries or all **application** entries for a specified physical disk.

The Delete operation queries you for the following information:

`Host Name`
: Select the host containing the physical disk for which you want to delete an application script.

`Disk Name`
: Select the physical disk module or modules for which you want to delete an application script.

`Application`
: Choose the application for which you want to delete entries.

Modifying an application Entry

To modify an entry in the **application** database, select the `Availability -> Disk Failover -> Application -> Modify` operation.

The Modify operation queries you for the following information:

`Host Name`
: Select the host containing the physical disk for which you want to modify an application script.

`Disk Name`
: Select the physical disk module or modules for which you want to modify an application script.

`Application`
: Select the application for which you want to modify entries.

`Remote Physical Diskname`
: Enter the name of physical disk on the remote system.

`Application Command Line`
: Enter the full pathname of a script or command to be executed when the local physical disk is failed over.

Listing application Entries

To display information about the entries in the **application** database, select the `Availability -> Disk Failover -> Application -> List` operation.

The entries in the **applications** database include the following information:

`Hostname`
: The remote host designated for failover.

`Local Diskname`
: The name of the physical disk as it appears on the current host.

`Remote Diskname`
: The name of the physical disk as it appears on the host designated for failover.

`Application Command`
: The full pathname of the script to be executed when the disk is failed over.

Managing the hosts Database

The **sysadm** `Availability -> Disk Failover -> Hosts` menu enables you to add, delete, and list entries in the **hosts** failover database. The following sections describe each operation.

Adding a hosts Entry

To add an entry to the **hosts** database, select the `Availability -> Disk Failover -> Hosts -> Add` operation. The name of the host to add is validated against the TCP/IP local and ONC/NIS **hosts** databases to ensure it is a valid host name. The Add operation queries you for the following information:

`Host name`
: Enter the name of the host to which a physical disk can be failed over. This must be a system that shares a disk in a dual-initiator configuration with your system.

Deleting a hosts Entry

To remove an entry in the **hosts** database, select the `Availability -> Disk Failover -> Hosts -> Delete` operation.

The Delete operation queries you for the following information:

`Host name`
: Select the host you want to delete.

Modifying a hosts Entry

To modify an entry in the **hosts** database, select the `Availability -> Disk Failover -> Hosts -> Modify` operation. This operation currently does nothing.

Listing hosts Entries

To display information about the entries in the **hosts** database, select the `Availability -> Disk Failover -> Hosts -> List` operation.

The entries in the **hosts** database include the following information:

`Hostname`
: The remote host designated for failover.

`Comms Path`
: The communication path to the host designated for failover.

`Remote Diskname`
: The current synchronization state of the **failover** databases for the designated host.

Giving Away a Failover Disk

To transfer control of a failover disk from your system to another, select the **sysadm** operation `Availability -> Disk Failover -> Give`. The operation requires that you enter the Host Name of the remote host and the Disk Name of the failover disks you wish to give. You should have already added the disks to the failover databases.

The **admfailoverdisk**(1M) command then consults the failover databases and performs the following actions:

- Terminates all user processes of the virtual disks and file systems on the specified physical disk.
- Unmounts the file systems and deregisters the physical disks.
- Registers the disks on the specified remote host.
- Checks and mount the file systems on the remote host.
- Executes any commands or scripts associated with the physical disk.

You cannot give *any* failover disk if your local giveaway database is out of sync, or marked NOTSYNC in the failover hosts file. Synchronize databases with the Synchronize operation, discussed later.

Taking Away a Failover Disk

To transfer control of a failover disk from another system to your own, select the **sysadm** operation `Availability -> Disk Failover -> Take`. The operation requires that you enter the host name of the remote host and the disk name of the failover disks you wish to take. You can take a failover disk only if

- The administrator of the remote system has added the disk to the failover databases, and
- The failover databases on your system have been successfully synchronized with the databases on the remote system.

The process of taking a failover disk involves having the **admfailoverdisk**(1M) code consult the failover databases and perform actions similar to those described previously for giving away a failover disk.

Optionally, you may select the Use Trespass option. This option is intended for situations where the remote system is hung or crashed. If the remote system is functioning normally, do not select the Use Trespass option. If you select the Use Trespass option when taking a failover disk from a normally functioning system, access to the failover disk on the remote system will be terminated ungracefully, possibly resulting in unnecessary data loss.

You cannot take a failover disk if your local takeaway database is out of sync, or NOTSYNC. Synchronize databases with the Synchronize operation, discussed later.

Verifying the Failover Database

To make sure that the local giveaway database has up-to-date information about failover disks and their virtual disks, or to list physical disk dependencies, select the operation `Availability -> Disk Failover -> Check`.

The operation requires that you specify the Host name of the remote system, the Disk name of the failover disk, and File name of the database files whose entries you would like to verify. The two files are the **giveaway** database, which contains information about physical and virtual disks designated to be given to other systems, and the **takeaway** database, which contains entries for physical and virtual disks designated to be taken from the other systems.

The operation scans the physical disk for all virtual disk pieces and correlates them with their entries in the local file system table, **fstab**, which you can display with the `File System -> Local Filesys -> List` operation. The operation then updates entries as necessary in the selected failover databases.

The operation also scans the physical disks for dependencies on other physical disks. A physical disk is considered dependent when it contains parts of virtual disks or mirrors that have parts on other physical disks. The operation reports any such dependencies that it finds. You may then add these other physical disks, using the `Availability -> Disk Failover -> Giveaway -> Add` operation, as necessary.

Synchronizing Failover Databases

To make the failover databases on your system consistent with those on a remote system, select the operation `Availability -> Disk Failover -> Synchronize`. For the specified remote system, the operation scans the local **giveaway** database and copies to the remote **takeaway** database entries for any virtual disks designated to be given to the remote system. The operation does *not* copy the remote takeaway disk entries to the local **giveaway** database.

Some operations that change the local **giveaway** database, such as Add, Delete, and Modify, offer the option of synchronizing databases. If you select this, **sysadm** copies the changes as appropriate to the takeaway database on the remote system. If you do not synchronize databases, the system flags the local **giveaway** database as out of sync, or NOTSYNC, and restricts you from giving any failover disks to another system until you have synchronized databases.

Machine Initiated Failover (MIF)

After your system is set up for OIF, you can set up Machine Initiated Failover (MIF) to eliminate the need for a human operator to detect a system failure and initiate the failover process. MIF has a monitoring process, **failovermon**(1M), that detects a failed host and initiates OIF disk takeover; the process can also detect a restored host and take steps to restore the disks taken over. The monitor process can use multiple communication paths, such as redundant LANs, to ensure that the monitored host has actually failed before taking action. When you define (add) a monitor process, you can specify that it is to run automatically on system reboot. If you do this, the process will run automatically at startup (by a command **sysadm** inserts in an init 3 script) until you stop the process or the system comes down.

MIF has two administrative commands, **admfailoveraltcommpath**(1M) and **failovermon**(1M), with associated **sysadm** interfaces. **admfailoveraltcommpath** is used to maintain the **altcommpath** database and to check alternate communication path accessibility. **failovermon** is used to maintain the **monitors** database and to perform failover monitoring. The manual pages provide detailed information.

MIF uses the following databases in addition to the OIF databases.

- **/etc/failover/altcommpath**. This contains information about alternate communication paths to the host that is being monitored. Entries contain the primary hostname of the host to be monitored and the hostname associated with the alternate communication path.
- **/etc/failover/monitors**. This contains information about the **failovermon** monitors that a system will run to monitor another system. Before another system can be monitored, there must be an entry for it in **monitors**.

Details on these two databases appear later in the chapter.

Setting Up and Using MIF

Before setting up MIF, you must set up and synchronize all of the OIF databases as explained earlier. Then follow these steps on each monitoring host to set up MIF. More details on each step appear in the following sections.

1. **Execute** `Availability -> Disk Failover -> Alternate Paths -> Add` on each monitoring host to set up the **altcommpath** database. Add the alternate communications paths that you want the monitor to use. **Sysadm** will request, and you must provide, the following values.

   - Primary host name: the network interface name of the host to monitor. DG/UX uses this name to look up alternate communication path entries.

   - Alternate remote host name: the secondary network interface name of the host to monitor. DG/UX uses this name to establish an alternate communication route to the host.

     After providing both names, tell **sysadm** to check the paths you specified; then confirm to perform the operation.

2. **Execute** `Availability -> Disk Failover -> Monitors -> Add` on the monitoring host to set up the **monitors** database and start the monitor.

   Sysadm will request, and you must provide, the following values.

   - Name of the host to monitor.

   - Seconds between heartbeats: the interval in seconds the monitor will sleep between message cycles.

   - Number of retries: the number of times the monitor will retry communication on each path before deciding the monitored host has failed.

   - Full pathname to the lost-pulse script to be executed when the monitoring host decides that the monitored host has failed. Data General supplies a default lost-pulse script in **/etc/failover/failovermon_lost_pulse**. You can edit this and use it or create another one.

   - Full pathname to the regain-pulse script to be executed when the monitoring host decides that the failed host has returned. Data General supplies a default regain-pulse script in **/etc/failover/failovermon_regain_pulse**.

   - **Yes** or **no** value indicating whether you want to start the monitor automatically on each system reboot.

   - **Yes** or **no** value indicating whether you want to start the monitor immediately after the successful completion of the Add operation.

The Monitoring Process

The failover monitor, **failovermon**(1M), is a daemon process that continues to run until you explicitly stop it or the system comes down. The monitor executes a loop of the following 6 tasks, shown in Figure 8–6. When you define (add) a monitor process, you can specify that it is to run automatically on system reboot. If you do this, the process will run automatically at startup (by a command **sysadm** inserts in an init 3 script) until you stop the process or the system comes down.

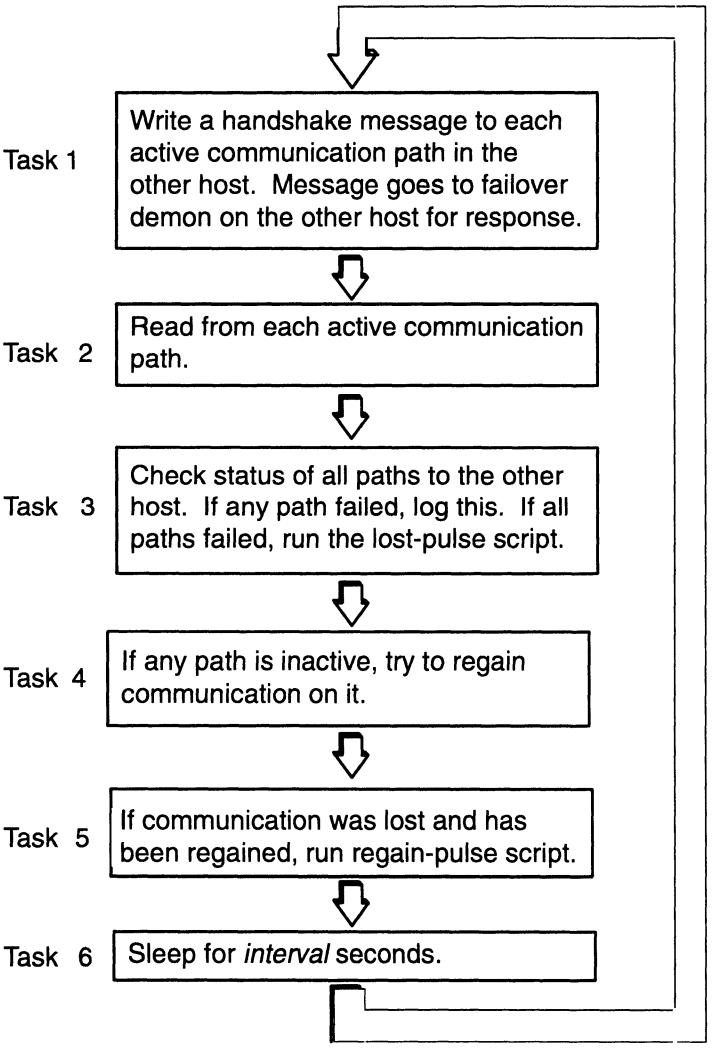

Figure 8–6 Tasks Performed by the MIF Failover Monitor Process

Managing the Alternate Paths Database (altcommpath)

The **sysadm** sequence `Availability -> Disk Failover -> Alternate Paths` lets you add, check, delete, modify, and list alternate paths. The following sections describe these operations.

Adding an Alternate Path

To add an alternate path for failover, use the **sysadm** sequence `Availability -> Disk Failover -> Alternate Paths -> Add`. **Sysadm** prompts for the following information:

`Primary Host Name [moe]`

Enter the primary host name of the remote host. If the remote host does not have an alternate hostname, the primary hostname is the only name.

`Alternate Remote Host Name [moe]`

Enter the alternate host name of the remote host. If the remote host does not have an alternate hostname, press Enter.

`Check path [no]`

Answer **yes** if you want to check the alternate communication path at the conclusion of the Add operation; otherwise answer **no**. Generally, we recommend a **yes** answer.

As with many other **sysadm** sequences, it asks for confirmation by prompting `Ok to perform operation`. Answer **yes** to add the alternate path as specified or **no** if you want to change a value.

If you told **sysadm** to check the path, it does so. The path check takes a few moments. If the alternate path is usable, **sysadm** will display

`Alternate path` *[alternate-remote-hostname]* `to` [*primary-hostname*] `is accessible.`

If the path is not usable, **sysadm** will display a message to that effect. Make sure the network(s) are operational; if they are, you should respecify the alternate path with Modify.

Checking an Alternate Path

To check an alternate path, use the **sysadm** sequence `Availability -> Disk Failover -> Alternate Paths -> Check`. **Sysadm** prompts for the primary and secondary host names as with Add above.

The path check takes a few moments. If the alternate path is usable, **sysadm** will display

`Alternate path` *[alternate-remote-hostname]* `to` [*primary-hostname*] `is accessible.`

If the path is not usable, **sysadm** will display a message to that effect. Make sure the network(s) are operational; if they are, you should respecify the alternate path with Modify.

Deleting an Alternate Path

To delete an alternate path, use the **sysadm** sequence `Availability -> Disk Failover -> Alternate Paths -> Delete`. **Sysadm** prompts for the primary and secondary host names as with Add above. After you confirm, it deletes the alternate path.

Modifying an Alternate Path

To change an alternate path, use the **sysadm** sequence `Availability -> Disk Failover -> Alternate Paths -> Modify`. **Sysadm** prompts for the primary and secondary host names and offers to check the path as with Add above. Enter the changed values you want. After you confirm, **sysadm** modifies the alternate path.

Listing Alternate Paths

To list the alternate paths, use the **sysadm** sequence `Availability -> Disk Failover -> Alternate Paths -> List`. **Sysadm** displays path information as in the following.

```
Hostname    Alt Hostname
--------    ------------

primary-    alternate-remote-
hostname    hostname
...         ...
```

Managing the Monitors Database

The **sysadm** sequence `Availability -> Disk Failover -> Monitors` lets you add, delete, modify, list, start, and stop monitors. The following sections describe these operations.

Adding a Monitors Entry

To add an entry to the **monitors** database, use the **sysadm** sequence `Availability -> Disk Failover -> Monitors -> Add`. **Sysadm** prompts for the following information:

`Host to Monitor`
: Specify the hostname you want this host to monitor for failure.

`Seconds Between Heartbeats: [0]`
: Enter the interval in seconds the monitor will sleep between message cycles to the other host. For example, **15**.

`Number of Retries: [0]`
: Enter the number of times the monitor will retry communication on each path before deciding the monitored host has failed. For example, **2**.

`Lost Pulse Action: [/etc/failover/failovermon_lost_pulse]`
: Enter the full pathname to the lost-pulse script to be executed when the monitoring host decides that the monitored host has failed. This lost-pulse script (and the companion regain-pulse script) is not the same as the application script specified to OIF. The lost-pulse and regain-pulse scripts are executed by the failover monitor; their main purpose is to include **sysadm** take commands that transfer control of the disks. The application script is executed by OIF after the disk transfer; its main purpose is to include commands that start application(s). Data General supplies a default lost-pulse script in **/etc/failover/failovermon_lost_pulse**. You can edit this and use it or create another one.

`Regain Pulse Action: [/etc/failover/failovermon_regain_pulse]`
: Enter the full pathname to the regain-pulse script to be executed when the monitoring host decides that the failed host has returned to service. See the comments under `Lost pulse action` above. Data General supplies a default regain-pulse script in **/etc/failover/failovermon_regain_pulse**. As with the lost–pulse script, you can edit this and use it or create another one.

`Start Monitor On Reboot? [no]`
: Specify the value **yes** or **no**, to indicate whether you want to start the monitor automatically on each system reboot. For example, **y**.

`Start Monitor Now? [no]`
: Specify the value **yes** or **no**, to indicate whether you want to start the monitor immediately after the successful completion of the Add operation. For example, **Y**.

As with many other **sysadm** sequences, it asks for confirmation by prompting `Ok to perform operation`. Answer **yes** to add the monitor as specified or **no** if you want to change a value.

Deleting a Monitor

To remove a monitored host from the **monitors** database, use the **sysadm** sequence `Availability -> Disk Failover -> Monitors -> Delete`. **Sysadm** prompts for the following information:

`Monitored Host: [`*xxx*`]`

Enter the hostname you want this host to stop monitoring. This hostname must have been added earlier to the **monitors** database.

Modifying a Monitor

To change the settings for a monitor entry in the **monitors** database, use the **sysadm** sequence `Availability -> Disk Failover -> Monitors -> Modify`. The **sysadm** prompts are exactly the same as for "Adding a monitors entry," above. Make any changes you want.

Listing Monitors

To list the entries in the **monitors** database, use the **sysadm** sequence `Availability -> Disk Failover -> Monitors -> List`. **Sysadm** displays monitor information as in the following.

```
Hostname  Act  Reb Int Ret  Lost Pulse                Regain Pulse
--------- --- --- --- ---  ------------------------  -------------------------

junior    yes  yes 15  2  /etc/failover/lost_pulse  /etc/failover/regain_pulse
```

The columns correspond roughly to the Add and Modify prompts. Act indicates active or not active status; Reb indicates whether or not the monitor will run on reboot; Int is the heartbeat interval; Ret is the number of retries, and Lost pulse and Regain pulse show the pathnames of the lost-pulse and regain-pulse scripts.

Starting a Monitor Process

To start a monitor process, use the **sysadm** sequence `Availability -> Disk Failover -> Monitors -> Start`. You must have already defined the monitor with the **sysadm** Monitors Add sequence. **Sysadm** prompts

`Host to monitor: [`*xxx*`]`

Enter the hostname you want this host to monitor. This hostname must have been added earlier to the **monitors** database.

After you answer and confirm, **sysadm** tries to start the monitor process on your system.

Stopping a Monitor Process

To stop a monitor process, use the **sysadm** sequence `Availability -> Disk Failover -> Monitors -> Stop`. The monitor process you want to stop must be running (usually started using a **sysadm** Monitor Add or Start sequence). **Sysadm** prompts

`Host to monitor: [`*xxx*`]`

Enter the hostname you want this host to stop monitoring.

After you answer and confirm, **sysadm** tries to stop the monitor process on your system.

Failover Example

This example shows how you might set up OIF and then MIF between two systems using shared buses. The systems run an application for Acme Products, S.A. The primary host, hostname **production**, runs the application; the backup host, hostname **development**, normally runs applications other than Acme but it monitors host **production** and stands ready to take over **production**'s disks if **production** fails.

The hardware configuration for the example is shown in Figure 8–7.

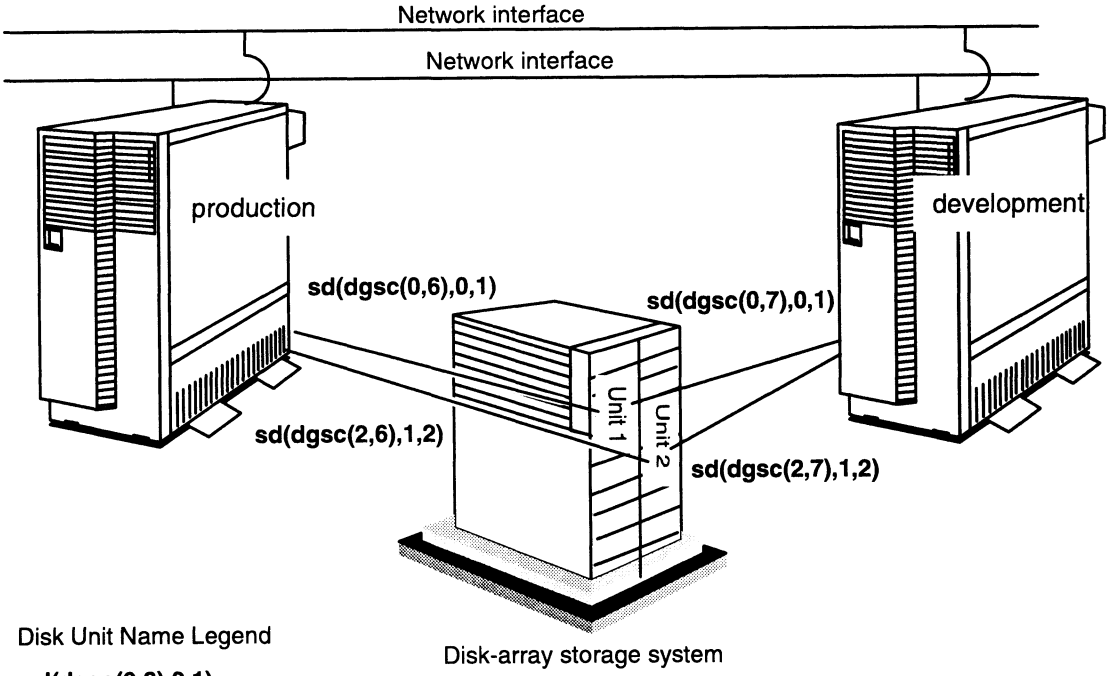

Physical Disk Layout

Route to disk from host production	Route to disk from host development	Virtual disk name and mount point directory	Application start-up and lost pulse scripts
sd(dgsc(0,6),0,1)	sd(dgsc(2,7),1,1)	acme, /acme	/acme/failover/acme/acme.up
sd(dgsc(0,6),0,1)	sd(dgsc(2,7),1,1)	acme_data (piece 1 of 2) /acme/data	/etc/failover/ failovermon_lost_pulse
sd(dgsc(0,6),0,2)	sd(dgsc(2,7),1,2)	acme_data (piece 2 of 2) /acme/data	None

Figure 8–7 Sample System and Physical Disk Layout with Two Shared Buses

In the interest of higher availability, the example shows *two* shared SCSI buses. This is also known as a dual-initiator-each-with-dual-adapter configuration. The disk-array storage system includes several virtual disks, including the Acme application virtual disk that spans two physical disks. The application script is executed by OIF after the disk transfer and includes commands needed to start the application.

Only one network interface is required, but again in the interest of higher availability, the example shows two interfaces. The hostnames on one network interface are **production** and **development**; on the other interface they are **production–alt** and **development–alt**.

Assume that DG/UX 5.4 release 3.00 is installed on the system disks, disks in the disk-array storage system are bound and soft formatted, DG/UX kernels for each host are built correctly, and the application and its data are loaded. OIF has not yet been set up.

Setting up OIF for the Example

1. Set up OIF. On **production**, you can use the following **sysadm** menu sequences:

   Execute `Availability -> Disk Failover -> Giveaway -> Add` and then specify **sd(dgsc(0,7),0,1)** (local physical disk); **sd(dgsc(0,6),0,1)** (remote physical disk); **development** (hostname to fail disk over to); **no** (to not synchronize, since you want to wait until you've added the next disk); Enter (to omit an application command line); and Enter (to confirm).

   For the next disk, execute Add and then specify **sd(dgsc(0,7),0,2)** (local physical disk); **sd(dgsc(0,6),0,2)** (remote physical disk); **development** (hostname to fail disk over to); **yes** (to synchronize); Enter (to omit an application command line); and Enter (to confirm).

2. Decide what actions should be taken in the application script whose pathname you specified above; then create the file in the specified directory and enter text in it. Generally, this script will do whatever is needed to start the application. The pathname is shown in Figure 8–7 as **/acme/failover/acme/acme.up**.

Setting Up MIF for the Example

Setting up MIF requires that the hardware, including network interfaces and disks, and software, including DG/UX kernels, be installed and configured. It also requires that OIF be set up.

Given these conditions, you can set up MIF by setting up alternate communications paths and the failover monitor. With OIF set up in the configuration shown above, follow these steps to implement MIF on the system **development**.

1. Set up the alternate communication path. You can use the **sysadm** menu sequence `Availability -> Disk Failover -> Alternate Paths -> Add` and then specify **production**, **production-alt**, **y** (to confirm the path) and **y** (to perform the operation).

2. Add the failover monitor and start it. You can use the **sysadm** menu sequence `Availability -> Disk Failover -> Monitors -> Add` and then specify **development**, **15** (seconds between heartbeats), **1** (retries), Enter (for default script **/etc/failover/failovermon_lost_pulse**), Enter (for default script **/etc/failover/failovermon_regain__pulse**), **yes** (to start monitor automatically on reboot), **yes** (to start monitor immediately) and Enter (to confirm).

3. Decide what actions should be taken in the **failovermon_lost_pulse** and **failovermon_regain_pulse** scripts. Since the application script starts the Acme application, we will not start the Acme software with the lost-pulse script. However, we can and will use the regain-pulse script to shut the Acme software down gracefully.

   Here is our **failovermon_lost_pulse** script.

```
#!/bin/sh
#
#  failovermon_lost_pulse shell script
#
#  This script uses a take with trespass to move the disks
#  used for the ACME software to the system development. This
#  script is executed when failovermon(1M) has detected that
#  host production has failed.
#
DISKLIST="sd(dgsc(0,7),0,1) sd(dgsc(0,7),0,2)"
#
admpdisk -o list -r "$DISKLIST" > /dev/null 2>&1
if [ $? -ne 0 ]
then admfailoverdisk -o take -T -h production "$DISKLIST"
else
echo "Disks $DISKLIST are already registered." >/dev/console
fi
## end of script
```

Here is the **failovermon_regain_pulse** script.

```
#!/bin/sh
#
# failovermon_regain_pulse shell script
#
# This script is run when host development detects that host
# production has come back on line. It takes down the ACME
# software and returns the physical disks for ACME to host
# production. We sleep at the beginning to ensure that
# production has finished running its /etc/rc3.d script for
# failover.
#
DISKLIST="sd(dgsc(0,7),0,1) sd(dgsc(0,7),0,2)"
sleep 30
/etc/failover/acme/acme_stop
admfailoverdisk -o give -h production "$DISKLIST"
# end of script
```

Disk Failover Troubleshooting

While testing failover, you can use two diagnostic aids: file **/var/adm/messages** and the **admpdisk** command or equivalent **sysadm** menu sequence.

Using /var/adm/messages

While testing a MIF setup, run **tail –f** on file **/var/adm/messages**. The failovermon writes to **/var/adm/messages**. If you watch this file while simulating a failure, you will see messages indicating that the communications path is inaccessible, and that the various action scripts are being run. If one of the scripts is not executable or not available, then information saying this will be written to **/var/adm/messages**.

admpdisk(1M)

The **admpdisk** command with various options will tell you whether a disk is registered or if it can be registered.

Managing IP (Internet Protocol) Takeover

This section builds on the previous major section "Managing Disk Failover" and assumes you understand that section.

The IP (internet protocol) takeover feature extends disk failover features to make a failed host's disk and resources available to NFS clients. IP takeover depends on and operates much the same way as Operator Initiated Failover (OIF) and Machine Initiated Failover (MIF). The features differ in that IP takeover transfers a floating network address from a host to another host, while disk failover transfers the disks themselves. Together the two features provide high availability of the failover disk file system to NFS client systems.

For routine operation, network hosts access the primary server host using this floating network address instead of a fixed, hard-wired address. If the primary host fails, the secondary host can take over the floating network address along with the failover disks. After a brief delay, NFS clients that have remote-mounted file systems on the failed host can continue running since the surviving host provides access to the file systems.

IP takeover is primarily useful because it provides higher availability of NFS remote-mounted file systems, telnet, and rlogin functionality. You can use IP takeover in conjunction with OIF or MIF. For the NFS clients to maintain access to the remote-mounted file system, the host's disk resources must have been transferred. You can transfer the disks manually using OIF or the DG/UX system can do it using an MIF failover script.

For IP takeover to work, two hosts on the same network or subnet must be connected in a configuration that will support OIF or MIF. That is, the hosts must be connected in a dual-initiator configuration (or with a disk-array storage system, a cabinet-sharing configuration).

Figure 8–8 shows two hosts with NFS clients in a model configuration for IP takeover.

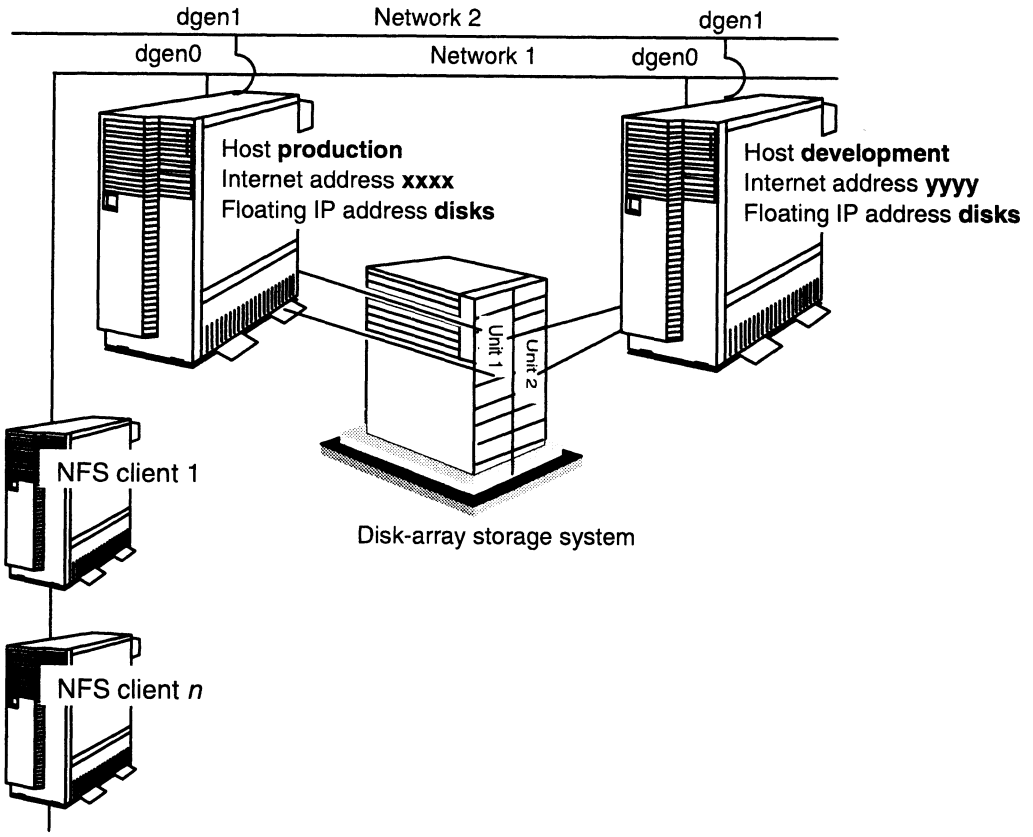

Figure 8–8 Two Hosts with NFS Clients in a Sample IP Takeover Configuration

In Figure 8–8, host **production** provides NFS services to all the NFS clients. If **production** fails, the clients will lose NFS services — until the floating IP address server is transferred (along with the pertinent disks in the disk-array storage system) from **production** to **development**). OIF/MIF provide the means to transfer the disks; IP takeover provides the means to transfer the IP address.

For continuity with the previous example, Figure 8–8 shows two network interfaces. The IP takeover example uses the network that runs between the **dgen0** interfaces only; **dgen1** serves as a secondary interface for the failover monitor.

Prerequisites for IP Takeover

IP takeover requires the following conditions to be true.

- The network between the hosts must use an Ethernet or FDDI LAN controller; token ring LANs are not supported. The LAN controllers may be of different types, such as VLCi and integrated.
- The two hosts must be on the same network or subnet.
- You must have added the hostname and floating IP address to the **/etc/hosts** database (using the **sysadm** sequence Networking –> TCP/IP –> Hosts), and to the NIS database (if applicable). The NFS clients must also have the hostnames and floating IP addresses added to these databases.

- The file systems to be involved in the takeover must be exported (made exportable and mounted via **sysadm**).
- You must have set up the network and host hardware and DG/UX software for OIF and, if you want automatic transfer of the network address, for MIF. Setting up OIF and MIF are explained earlier in this chapter.

You do not have to change the OIF/MIF hardware configuration or build a new kernel to use IP takeover.

Setting Up and Using IP Takeover

To set up IP takeover for two hosts, you must execute the following steps (assuming the prerequisite conditions above have been fulfilled).

1. On the host that will be the primary NFS server on the network, with the network running, add the IP takeover entry. To do this, use the **sysadm** sequence `Availability -> IP Takeover -> Add`. When you have added the IP entry, be sure to answer **yes** to the synchronize prompt. The synchronize step enters the floating IP address in the IP takeover database of the other host, allowing it to take over the name. If you do not synchronize databases, the system flags them as out of synchronization (noted as NOTSYNC on the **sysadm** `Availability -> Disk Failover -> Hosts -> List` **status** display); and it will not allow you to transfer control of the IP address.
2. With the IP takeover entry added and databases synchronized, you can transfer the IP address as necessary. There are several ways to transfer this address, depending on whether both hosts are running, as follows.
   - When the giving host has failed, you can use the taking host to transfer control to the taking system with the **sysadm** sequence `Availability -> IP Takeover -> Take` and the **Use Trespass** option.
   - When the giving host is functioning normally, you can transfer control using either system:

     from the giving host using the **sysadm** sequence `Availability -> IP Takeover -> Give`; or

     from the taking host using the sequence `Availability -> IP Takeover -> Take` *without* the **Use Trespass** option.

Managing the IP Takeover Database

On each host, DG/UX uses the database file named **/etc/failover/failoverip** to manage IP takeover. The **failoverip** database contains information about IP host names and the network interface used on each host. You maintain this database using a **sysadm** menu sequence described here or the equivalent **admfailoverip**(1M) command. A manual page provides detailed information on this command.

The **sysadm** sequence `Availability -> IP Takeover` lets you add, delete, modify, and list database entries; give and take the floating IP address; start and

stop the floating IP address mechanism; and synchronize databases. The following sections describe each operation.

Adding an IP Takeover Entry

Before anyone can transfer the floating IP address to another host, you must add the address to the IP takeover database using the `Availability -> IP Takeover -> Add` operation. After entering the needed information, you synchronize databases to add the information to the other host's failover databases.

The Add operation prompts for the following information:

`Host name:`
: Enter the name of the remote host that will get the IP address if your host fails. To continue the example from the disk failover section, host **production** would add hostname **development**). Figure 8–8, earlier, shows sample hosts.

`Floating Address Name:`
: Enter the name for the floating IP address. Unlike most hostnames, which are associated with a single interface on a single system, the floating address name can float between the two hosts for use by either host's network interface. This name will allow NFS clients to access remote file systems on disks that were transferred to the other host by OIF/MIF scripts. (The floating IP address must always belong to the same host that the disks belong to.) NFS clients will mount the remote file system using the floating address name you specify here; they will not use a hostname.

  Since the same floating address name will serve on both hosts, choose a name that identifies the disk resources, not a host. The floating IP address is unique to IP takeover and must differ from other hostnames on the network. You can use as many as nine DG/UX filename characters. For example, to continue the example from the disk failover section, **acmedisks**.

`Local Network Interface:`
: Enter the name of the local host's network interface to use for IP takeover. For example, **dgen0** or **cien0**.

`Remote Network Interface:`
: Enter the name of the remote host's network interface to use for IP takeover. For example, **dgen0** or **cien0**.

`Start Floating IP Address On System Reboot? [yes]`
: Specify **yes** or **no** to indicate whether you want to start the IP takeover mechanism automatically on each system reboot.

`Start Floating IP Address Now? [no]`
: Specify **yes** or **no** to indicate whether you want to start the IP takeover mechanism immediately after the successful completion of the Add operation.

`Synchronize database? [no]`

This option updates the information about this floating IP address entry in the failover databases on the remote system. If the remote system is inaccessible, synchronization will fail. When the remote system becomes accessible again, synchronize databases with the Synchronize operation, explained later. Until you synchronize databases, the system considers the local IP failover database as out of synchronization (noted as NOTSYNC on the **sysadm** `Availability -> Disk Failover -> Hosts -> List` status display). With the databases out of synchronization, no one can use the IP takeover mechanism with that host.

Adding the IP takeover entry is the essential software task required for IP takeover. For IP takeover maintenance tasks, continue with the next sections.

Deleting an IP Takeover Entry

To remove an IP takeover entry from the IP failover database, select the operation `Availability -> IP Takeover -> Delete`. **Sysadm** prompts for the following information:

`Host name:`

Enter the name of the remote host from whose IP failover database you want to delete the floating IP address.

`Floating Address Name:`

Enter the name of the floating IP address you want to delete.

`Synchronize database? [no]`

This option removes the information for the floating IP address from the other host's failover databases. We recommend that do this. For more information, see "Synchronize database" under the previous section "Adding an IP Takeover Entry."

Modifying an IP Takeover Entry

To change the information about an IP takeover entry in your host's failover database, select the operation `Availability -> IP Takeover -> Modify`. **Sysadm** prompts for the hostname, floating address name, and other information as shown in the earlier section "Adding an IP Takeover Entry."

Listing IP Takeover Entries

To display information about local IP takeover entries, select `Availability -> IP Takeover -> List`. The display includes the following information:

- the hostname of the remote host;
- the floating IP address name to share;
- the local network interface name;
- the remote network interface name;
- yes or no flag that shows whether the floating IP address mechanism will be restarted when the system is rebooted; and
- hostname the floating IP address is active on, or NONE if the address is not active on either host.

For synchronization status (not included on the List screen), use the **sysadm** sequence `Availability -> Disk Failover -> Hosts -> List`.

Giving a Floating IP Address

To transfer control of an IP Takeover address from your host to another, select the **sysadm** operation `Availability -> IP Takeover -> Give`. For transfer to work, the failover databases must be synchronized, as noted on the `Availability -> Disk Failover -> Hosts -> List` status display. If they are not synchronized (NOTSYNC), synchronize the databases with the Synchronize operation, explained later.

NOTE: Transferring a floating IP address prevents NFS clients from accessing your host by the network interface associated with this address. You should do this only after the failover disk(s) have been transferred by OIF or MIF operations, explained in the previous section on disk failover. NFS clients may regain access to your host through the other host after a few moments. Any user who was running a telnet or rlogin session must restart the session.

Sysadm prompts for the following information.

`Host Name:`
: Enter the name of the remote host that will get the IP address.

`Floating Address Name:`
: Enter the name for the floating IP address you want to transfer to the other host.

Sysadm then prompts for confirmation. After you confirm, **sysadm** consults the IP failover databases and performs the following actions:

- on the local host, checks the IP takeover database and, if the floating IP address is active on this host, deactivates the address; and
- on the remote host, activates the floating IP address.

Taking a Floating IP Address

The process of taking a floating IP failover address involves having the **admfailoverip**(1M) code consult the failover databases and perform actions similar to those described previously for giving away a floating IP address.

To transfer a floating IP address from another host to your own, select the **sysadm** operation `Availability -> IP Takeover -> Take`. The operation requires that you enter the remote hostname and the floating address name you wish to take. You can take a floating IP address only if the administrator of the remote system has added a floating IP address for your host to the failover databases and has synchronized databases.

Optionally, you may select the Use Trespass option. This option is intended for situations where the remote system is down and therefore unable to respond. If the

remote system is functioning normally and you select the Use Trespass option, the transfer will fail and your system will display an error message.

Starting the IP Takeover Mechanism

Normally, the IP mechanism starts automatically, according to the answer you gave to the start-on-reboot prompt when you added or modified this IP takeover entry. To start (enable) the IP Takeover mechanism on your host, use the **sysadm** sequence `Availability -> IP Takeover -> Start`. **Sysadm** prompts for the remote hostname for which you want to start IP takeover and then prompts for the floating address name. Starting of takeover affects only the current state on the local host; remote host databases are not affected by this. When the start operation succeeds, **sysadm** will display a "Floating IP address started" message.

Stopping the IP Takeover Mechanism

To stop (disable) the IP Takeover mechanism on your host, use the **sysadm** sequence `Availability -> IP Takeover -> Stop`. **Sysadm** prompts for the remote hostname for which you want to stop IP takeover and then prompts for the floating address name. Stopping of takeover affects only the current state on the local host; remote host databases are not affected by this. The action on the next reboot will depend on the answer you gave to the start-on-reboot prompt when you added or modified this entry.

Synchronizing Databases

To make all the failover databases on a remote system consistent with the ones on your system, select the operation `Availability -> IP Takeover -> Synchronize`. **Sysadm** then prompts for the hostname whose databases you want to synchronize with yours. **Sysadm** scans the local databases and copies to the remote host the entries for that host. **Sysadm** does *not* copy the remote **failoverip** entries to the local database.

Some operations that change the local IP takeover database, such as Add, Delete, and Modify, offer the option of synchronizing databases. If you select this option, **sysadm** copies the changes as appropriate to the **failoverip** database on the remote system. If you do not synchronize databases, the system flags the local databases as out of sync, or NOTSYNC, and restricts you from giving any floating IP address to another system until you have synchronized databases.

Mounting File Systems for Use with IP Takeover

For the NFS clients to access the file systems on the failover disks, the clients must mount the file systems using the floating IP network address. This allows the clients to access the file system as long as the floating IP network address is available.

For two serving hosts to be able to mount a file the original owning host must use the floating IP network address when it originally mounts the file system; if the host does not use the floating IP address, the physical disk failover may fail.

Make sure the mount point directory is what your application expects. If you must change the local mount point directory name to satisfy the application or IP failover, then you will need to change mount point specification in OIF giveaway database and then synchronize databases.

Technical Details on Floating IP Mounts

When one of the server hosts uses the floating IP address to mount a file system, this results in a two–level mount scheme on that server. The lower-level mount will be a file system on the actual virtual disk. It results in an **/etc/fstab** entry of the following form:

/dev/dsk/*virtual-disk-name* */file-system-name* **dg/ux rw x 0**

For example,

```
/dev/dsk/acme_data       /acme_data_fs   dg/ux rw x 0
```

This file system (*/file-system-name*) needs to be exported when the physical disk is added to the failover databases. The OIF software will store the export information in the failover databases and ensure the file system is mountable by NFS clients. This lower level mount will be managed by the OIF software during a failover of the disk. Because the file system is not directly accessed (no process opens a file in */file-system-name*), this mount can be managed when the system controlling the disks fails.

The higher level mount is to mount the exported file system the same way that the NFS clients. This results in a **/etc/fstab** entry of the following form:

floating-ip-address:*/file-system-name* */mount-point-directory* **nfs rw,hard,intr,bg x x**

For example,

```
acmedisks:/acme_data_fs   /acme   nfs  rw,hard,intr,bg  x  x
```

All applications that need to access the file system will do so by accessing */mount-point-directory*; in this example, **/acme**.

IP Takeover Example

Figure 8–9 shows and example of IP failover. This figure duplicates Figure 8–8 (for your convenience) and builds on the OIF and MIF example shown in the section on disk failover.

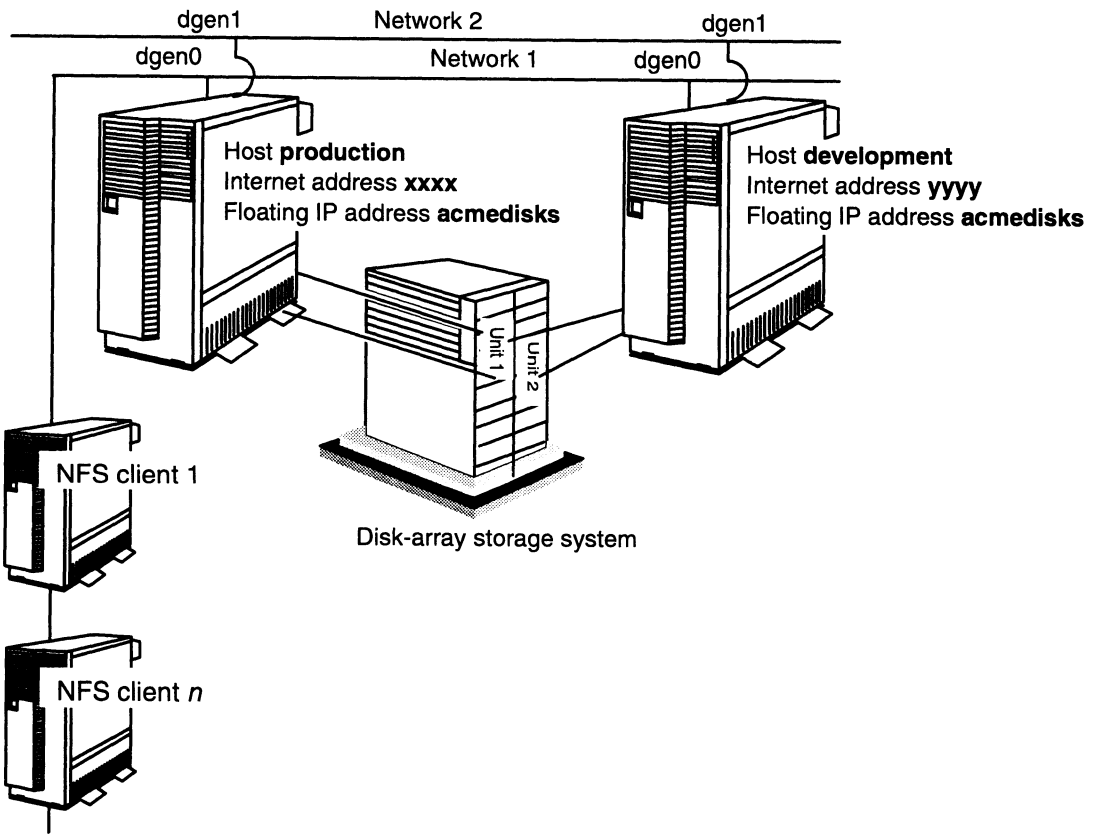

Figure 8–9 Two Hosts with NFS Clients in a Sample IP Takeover Configuration

As before, this shows two networks, but the IP takeover example uses the network that runs between the **dgen0** interfaces only; **dgen1** serves as a secondary interface for the failover monitor.

Setting up IP Takeover for the Example

1. On host **production**, for the first network interface, add the hostname **development** and other information as follows.

   Use **sysadm** sequence `Availability -> IP Takeover -> Add` and then specify the following information:

```
Host Name: [junior] development↵
Floating Address Name: acmedisks↵
Local Network Interface: dgen0↵
Remote Network Interface: dgen0
Start Floating IP Address on System Reboot? [yes] ↵
Start Floating IP Address Now? [no] yes↵
Synchronize database? [no] yes↵
OK to perform operation? [yes] ↵
```

2. Decide what IP takeover action you want the backup host, **development**, to take if the primary host **production** fails and when **production** resumes normal operation. You will need to insert the appropriate **sysadm admfailoverip** command in the MIF failover monitor **lost_pulse** and **regain_pulse** scripts on **development** shown earlier. In each script, the **admfailoverip** command must follow the command that transfers the disk.

   Following are two examples of the **lost_pulse** and **regain_pulse** scripts, based on those in the previous section. The first script, on **development**, is **failovermon_lost_pulse**. The lines added for IP takeover are near the end.

```
#!/bin/sh
#
#  failovermon_lost_pulse shell script
#
#  This script uses a take with trespass to move the disks
#  used for the ACME software to the system development. This
#  script executes when failovermon(1M) has detected that
#  host production has failed.
#
DISKLIST="sd(dgsc(0,7),0,1) sd(dgsc(0,7),0,2)"
#
admpdisk -o list -r "$DISKLIST" > /dev/null 2>&1
if [ $? -ne 0 ]
then admfailoverdisk -o take -T -h production "$DISKLIST"
else
echo "Disks $DISKLIST are already registered." >/dev/console
fi
# For IP takeover, we added this and the following three lines
# to the OIF/MIF script:
# Use IP takeover to take over the floating IP address.
admfailoverip -o  take -h  production  -T  acmedisks

# end of script
```

The **failovermon_regain_pulse** script, still on **development**, follows. Again, the lines added for IP takeover are near the end.

```
#!/bin/sh
#
# failovermon_regain_pulse shell script
#
# This script runs when host development detects that host
# production has come back on line. It takes down the ACME
# software and returns the physical disks for ACME to host
# production. We wait (sleep) at the beginning to ensure that
# production has finished running its /etc/rc3.d script for
# failover.
#
DISKLIST="sd(dgsc(0,7),0,1) sd(dgsc(0,7),0,2)"
sleep 30
/etc/failover/acme/acme_stop
admfailoverdisk -o give -h production "$DISKLIST"
# For IP takeover, we added this and the following three lines
# to the OIF/MIF script:
# Use IP takeover to give the floating IP address.
admfailoverip -o  give -h production acmedisks
# end of script
```

3. Each NFS client on this network would add the **acmedisks** file system as a remote file system, using a sequence like the following.

```
File System  -> Remote Filesys -> Add

Mount Directory: /acme↵
Remote Host Name: acmedisks↵
Remote Mount Directory: /acme↵
Write Permission: [Read/Write]↵
NFS Mount Type: [Hard] ↵
Interruptible? [yes] ↵
Retry in background? [yes]↵
Mount the file system? [yes]↵
OK to perform operation? [yes]↵
```

Adding the remote file system on each NFS client mounts the file system and enters its name in file **/etc/fstab**, from which DG/UX will mount it automatically on future client startups.

Troubleshooting IP Takeover

The failover monitor writes to the file **/var/adm/messages**. While testing IP takeover, as with other types of failover, you can look at the file **/var/adm/messages** for messages.

While simulating a failure, run **tail –f** on file **/var/adm/messages**. You will see messages indicating that the communications path is inaccessible, and that the

various action scripts are being run. If one of the scripts is not executable or not available, then information saying this will be written to **/var/adm/messages**.

After starting the IP takeover mechanism, you can check status from the starting host as follows:

- To make sure that the floating IP address interface (in the example, **acmedisks**) has been added to the file **/etc/tcpip/params**, examine that file.
- To make sure that the failover database shows the exported file systems, on the host that has the disks, use the **sysadm** sequence `Availability -> Disk Failover -> Giveaway -> List`; on the host that does not have the disks, use the **sysadm** sequence `Availability -> Disk Failover -> Takeaway -> List`. The export flag field of the display should show "YES"and this should be followed by any export options specified.
- To make sure the file system has been exported, on the host that has the disks, use the sysadm sequence `File System -> Local Filesys -> List` and at the query "Which local file systems?," answer with "exported."
- To make sure that the interface is up and running, on the host that has the disks, type

```
# admipinterface  -o list  -g  failover ↵
```

From an NFS client, to make sure that the floating IP address interface is available, you can type

ping *floating-ip-address-name*

An example follows:

```
# ping acmedisks↵
```

Managing Multi-Path LAN I/O

The High Availability features discussed in the previous sections were for configurations with multiple systems. Multi-path LAN I/O is a feature that enables you to protect a single system from LAN interface failures.

With multi-path LAN I/O, you can configure your system to automatically reroute LAN I/O. You can configure most types of LAN interfaces in AViiON systems so that two LAN interface cards can be connected to the same LAN. If the primary LAN interface fails, you can set the DG/UX system to automatically switch from the primary LAN interface to a backup LAN interface. The process is transparent to applications and protects your system from both LAN interface failures and I/O channel failures.

Multi-path LAN I/O has the following restrictions:

- Your system must support two LAN network interfaces, one of which is idle
- Both interfaces must be the same type

- LAN network interfaces must be one of the following types:
  - **cien**
  - **dgen**
  - **hken**
  - **pefn**

You must add the **mpl()** device driver to the kernel to implement multi-path LAN I/O. Note that the driver is not added to the kernel by default. You must add it explicitly, rebuild your kernel, and reboot your system. Once the **mpl** driver is in the kernel (that is, **/dev/mpl** is present), you can use **sysadm** to set up your two LAN interfaces in a multi-path relationship. **Sysadm** adds an entry to the **/etc/iopath.params** database describing the multi-path relationship.

Following is an example of the **iopath.params** file:

```
#########################################################################
#
#                       iopath database parameters
#
#########################################################################
#
#  The entries in this file correspond to the two I/O path names
#  that will be bound together to allow the system to perform
#  multi-path I/O. The format of these entries is:
#
#       <primary_path>          Name of primary I/O path.
#       <start_on_reboot>       Flag indicating relationship starts
#                               when system is rebooted.
#       <backup_path>           Name of first backup I/O path.
#       <backup_path>           Name of second backup I/O path.
#
#------------------------------------------------------------------------
dgen0 YES dgen1 NONE
```

The name of the primary I/O path is used as the key in all related **sysadm** operations to locate entries in this database. For more information on the **iopath.params** database, see the **iopath.params**(4M) manual page.

The multi-path LAN I/O process follows this sequence of events:

1. The **failovermon** monitor periodically requests that the **mpl** driver have both LAN interfaces send a message intended for the other interface.

   For more information on **failovermon**, see the **failovermon**(1M) manual page.

2. If both interfaces receive their message, the **mpl** driver determines that the I/O paths are open and takes no further action until the next check.

3. If only one interface receives a message, the **mpl** driver uses that controller's driver to determine the state of the interface.

4. If that interface has failed, the **mpl** driver switches I/O to the remaining interface.

Adding Multi-Path LAN I/O Entries

Select `Availability -> Multi Path I/O -> LAN -> Add` to add an entry to the **iopath.params** database. The Add operation displays the following queries:

`Primary LAN I/O Path`

Enter the device specification of the LAN interface to be used as the primary I/O path. For example, you could enter `dgen0`.

The primary path should be the interface used by the upper level networking software, such as TCP/IP. In the case of TCP/IP, the primary interface should be the one listed in the **/etc/tcpip.params** file.

`Backup LAN I/O Path`

Enter the device specification of the secondary LAN interface to be used as the backup I/O path. For example, you could enter `dgen1`. Note that this controller must be of the same type as the primary interface. For example, both controllers could be of type **dgen**.

The backup path should not be the interface used by default by the upper level networking software. If the backup interface is being used, Multi-path LAN I/O will not start for that pair of interfaces.

`Start Multi-Path LAN I/O on Reboot?`

Enter `yes` if you want to automatically start this multi-path relationship when the system boots. Enter `no` if you do not. This process is controlled by the **rc.failover** script.

`Start Multi-Path LAN I/O Now?`

Enter `yes` if you want to start this multi-path relationship when you complete the Add operation. Enter `no` if you do not.

If a **failovermon** monitor is not presently running on the system, this operation automatically starts the monitor to control the timing of the **mpl** driver's checks on the state of the I/O paths.

Deleting Multi-Path LAN I/O Entries

Select `Availability -> Multi Path I/O -> LAN -> Delete` to stop a multi-path LAN I/O relationship and delete its entry from the **iopath.params** database. The Delete operation displays the following query:

`Primary LAN I/O Path`

Enter the device specification of the primary LAN interface in the multi-path LAN I/O relationship you want to delete. For example, you could enter `dgen0`.

This operation stops the multi-path I/O relationship for the specified controller and deletes its entry from the **iopath.params** database.

Modifying Multi-Path LAN I/O Entries

Select `Availability -> Multi Path I/O -> LAN -> Modify` to change an entry in the **iopath.params** database. The Modify operation displays the same queries as the Add operation:

`Primary LAN I/O Path`
: Enter the device specification of the LAN interface to be used as the primary I/O path.

`Backup LAN I/O Path`
: Enter the device specification of the secondary LAN interface to be used as the backup I/O path.

`Start Multi-Path LAN I/O on Reboot?`
: Enter `yes` if you want to automatically start this multi-path relationship when the system boots. Enter `no` if you do not.

`Start Multi-Path LAN I/O Now?`
: Enter `yes` if you want to stop and restart the modified multi-path relationship when you complete the Modify operation. Enter `no` if you do not.

Displaying Multi-Path LAN I/O Entries

Select `Availability -> Multi Path I/O -> LAN -> List` to print an **iopath.params** database entry to standard output. The List operation displays the following query:

`Primary LAN I/O Path`
: Enter the device specification of the primary LAN controller in the **iopath.params** database entry you want to list. For example, you could enter `dgen0`. Entering `all` lists all entries in the database.

The List operation reports the following information:

- Name of the primary I/O path
- Flag indicating whether to start the multi-path relationship when the system boots
- Name of the backup I/O path

  The second backup path is not used in multi-path LAN I/O.
- Name of the currently active I/O path

Following is an example of the List operation's output:

```
Primary Path       Rbt  Backup Path       Backup Path       Active Path
------------       ---- -----------       -----------       -----------
dgen0               YES  dgen1             NONE               dgen0
```

Starting Multi-Path LAN I/O

Select `Availability -> Multi Path I/O -> LAN -> Start` to start the multi-path LAN I/O relationship for a specified I/O path. The Start operation displays the following query:

`Primary LAN I/O Path`

Enter the device specification of the primary LAN interface in the multi-path LAN I/O relationship you want to start. For example, you could enter `dgen0`. Entering `all` starts all the relationships listed in the **iopath.params** database.

If a **failovermon** monitor is not presently running on the system, this operation automatically starts the monitor to control the timing of the **mpl** driver's checks on the state of the I/O paths.

Stopping Multi-Path LAN I/O

Select `Availability -> Multi Path I/O -> LAN -> Stop` to stop the multi-path LAN I/O relationship for a specified I/O path. The Stop operation displays the following query:

`Primary LAN I/O Path`

Enter the device specification of the primary LAN interface in the multi-path LAN I/O relationship you want to stop. For example, you could enter `dgen0`. Entering `all` stops all the relationships listed in the **iopath.params** database.

Note that if the backup LAN interface is the currently active I/O path, this operation does not work.

Switching to an Inactive LAN I/O Path

Select `Availability -> Multi Path I/O -> LAN -> Switch` to manually switch LAN I/O from the currently active I/O path to an inactive path. The Switch operation displays the following queries:

`Primary LAN I/O Path`

Enter the device specification of the primary LAN interface in the multi-path LAN I/O relationship for which you want to switch I/O paths. For example, you could enter `dgen0`.

`Backup Physical LAN`

Enter the device specification of the inactive I/O path to which you want to switch LAN I/O. For example, you could enter `dgen1`.

Indicating that a LAN I/O Path is Repaired

Select `Availability -> Multi Path I/O -> LAN -> Repaired` to indicate that a failed path in a multi-path LAN I/O relationship has been repaired. An example of a situation when you would need to use this operation is when the cable to a controller for a **hken** interface becomes detached. In this case, multi-path LAN I/O would automatically switch to the backup interface. However, after you reconnect the cable, you would need to use the Repaired operation on the primary interface before you could switch I/O back to it.

The Repaired operation displays the following queries:

`Primary LAN I/O Path`
: Enter the device specification of the primary LAN interface in the multi-path LAN I/O relationship for which you want to indicate an I/O path has been repaired. For example, you could enter `dgen0`.

`Backup Physical LAN`
: Enter the device specification of the repaired I/O path.

You should use this operation only after a failed interface has been repaired or replaced and then reconfigured into your system. However, if you had to bring your system down to repair the problem, you do not need to use the Repaired operation.

Note that this operation does not change the currently active I/O path. It just tells the system that a failed path has been repaired and is now available for use.

End of Chapter

Chapter 9
Managing File Systems

File system management tasks apply to all systems, even to those that do not have their own disks. File system management involves creating file systems and making them available on the system. It also involves verifying file system consistency, backing up and restoring files, and maintaining the file system table, **/etc/fstab**. The file system table contains entries for the local and remote file systems that you want available on your system.

File System Terms

You may also want to review the information in Chapter 7, which discusses virtual disks, physical disks, and features for improving disk and file system service.

We use the following terms in this chapter:

file system
A file system is a software–formatted portion of disk space typically located on a virtual disk. The file system contains the internal data structures that the operating system requires to keep track of files and directories on the virtual disk. Typically, you build a file system on each virtual disk that you create except on virtual disks used as swap area (for demand paging) and any virtual disks to be used for direct access by applications such as database managers. A file system typically is placed directly on a diskette, however, without an underlying virtual disk. See Chapter 15 for more information on diskette preparation.

/etc/fstab The **fstab**(4) file is the file system table describing both local and remote file systems that you want accessible on the local system and swap areas. File systems listed in **fstab** become accessible automatically at boot time. The **fstab** file contains information for commands that mount, unmount, backup, restore, and check file systems.

mount point directory
After you create a virtual disk, you usually create a file system on it. Next you mount the file system on a directory. From then on, the name of the file system is the name of the directory where you mounted it, called the mount point directory. When operations such as Backup, Restore, Check, or Disk Use require the name of a file system, you provide the name of the mount point directory.

mount To attach a file system to a directory, making it accessible to users. The system mounts directories listed in **fstab** at boot time. You can mount directories explicitly with the shell command **mount**(1M) or with **sysadm**'s Mount operation. You mount local file systems as well as remote ones. To mount remote file systems, you need the ONC/NFS network software.

unmount To detach a file system from a directory, making it inaccessible to users. The system unmounts remote (ONC/NFS-mounted) file systems when you shut your system down to run level 2 or lower. The system unmounts local file systems (except **/usr** and /) when you shut the system down to run level S (single-user mode).

backup cycle list
A plan for doing daily, weekly, and monthly backups on tape. A default backup cycle list is supplied with the DG/UX system.

pass number
A number from a file system's **fstab** entry that indicates the order in which the **fsck** program checks file systems for corrupted and damaged files. In the first pass, **fsck** simultaneously checks all file systems with a pass number of 1. In the second pass, **fsck** checks file systems with a pass number of 2, and so on.

read-write mode
Access permission that allows users to write (change) files and directories in the file system as well as read them. This mode does not override the normal permissions that already apply to the files and directories.

read-only mode
Access permission that allows only reading of a file system.

ONC/NFS Network File System. A network software package that allows you to access remote file systems as though the remote file systems resided on a local disk. For more information, see *Managing ONC™/NFS® and its Facilities on the DG/UX™ System.*

export To make a local file system available for mounting by other systems in your network. To export a local file system or mount an exported file system from a remote system, both systems must be running the ONC/NFS software.

fast recovery file system
A file system mounted with the **fsck** logging feature. The system records file system modifications to a log, which **fsck** uses during recovery after a failure. Using the log, **fsck** can check and repair the file system faster than it can without the log.

The Operating System's View of File Systems

From the operating system's perspective, with the exception of entire physical devices used for data storage, a file system typically is associated with a virtual disk, which in turn is associated with sections of physical disks accessed through a device node. For each device and virtual disk, the kernel creates a device node each time you boot the system. Virtual disks form the bridge between file systems and physical disks. The file system associated with the virtual disk is mounted (made available to users) in the directory structure.

You can think of the relationship between file systems, virtual disks, and physical disks as a three–level hierarchy.

File system The level at which the user interacts with the system. Optionally, you may create a file system when you create a virtual disk. Alternatively, you may create a file system on an existing virtual disk through operations for managing file systems (covered in this chapter).

Virtual disk The intermediate level that associates the file system with the physical disk. The file system and the rest of the operating system interact at this level. You create virtual disks by name.

Physical disk The level at which the operating system interacts with the hardware. Optionally, you may specify the physical disk on which the virtual disk is created.

Some virtual disks, such as the **swap** virtual disk or an add–on database manager package, do not require a file system. The operating system accesses those virtual disks directly without file system intervention.

Figure 9–1 shows the operating system's perspective toward a file system.

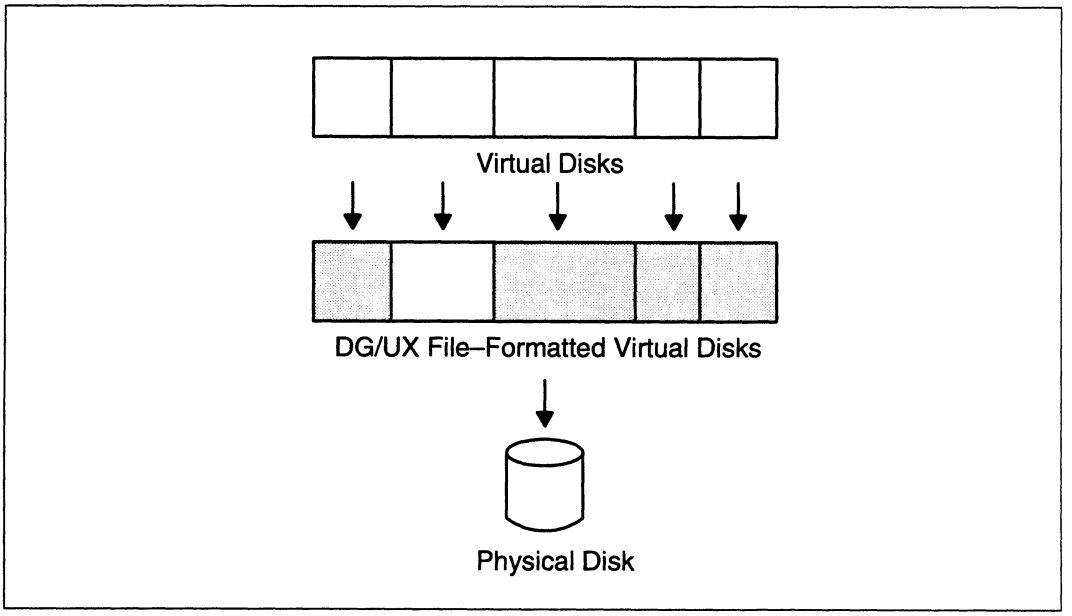

Figure 9–1 The Operating System's Perspective of a File System

The shaded virtual disks contain DG/UX file system structures. The unshaded one will contain software which does not require a file system; it will interact with the operating system directly.

Once a file system is superimposed on a virtual disk, you further characterize it with a file name and mount point location. Information on this procedure is given in this chapter.

The User's View of File Systems

From a user's viewpoint, there appears to be one file system: the single hierarchy consisting of all files and directories on the system. Because all file systems are

mounted under the / (root) file system, a system's entire directory tree appears to be a single hierarchical directory structure. If you mount a new file system, the directory tree (as seen by the user) expands, starting at the point where the new file system is mounted. The user sees a group of new files under a new directory name. If you unmount a file system, part of the directory tree disappears from view, but is inaccessible.

Figure 9–2 shows the user's perspective toward a file system.

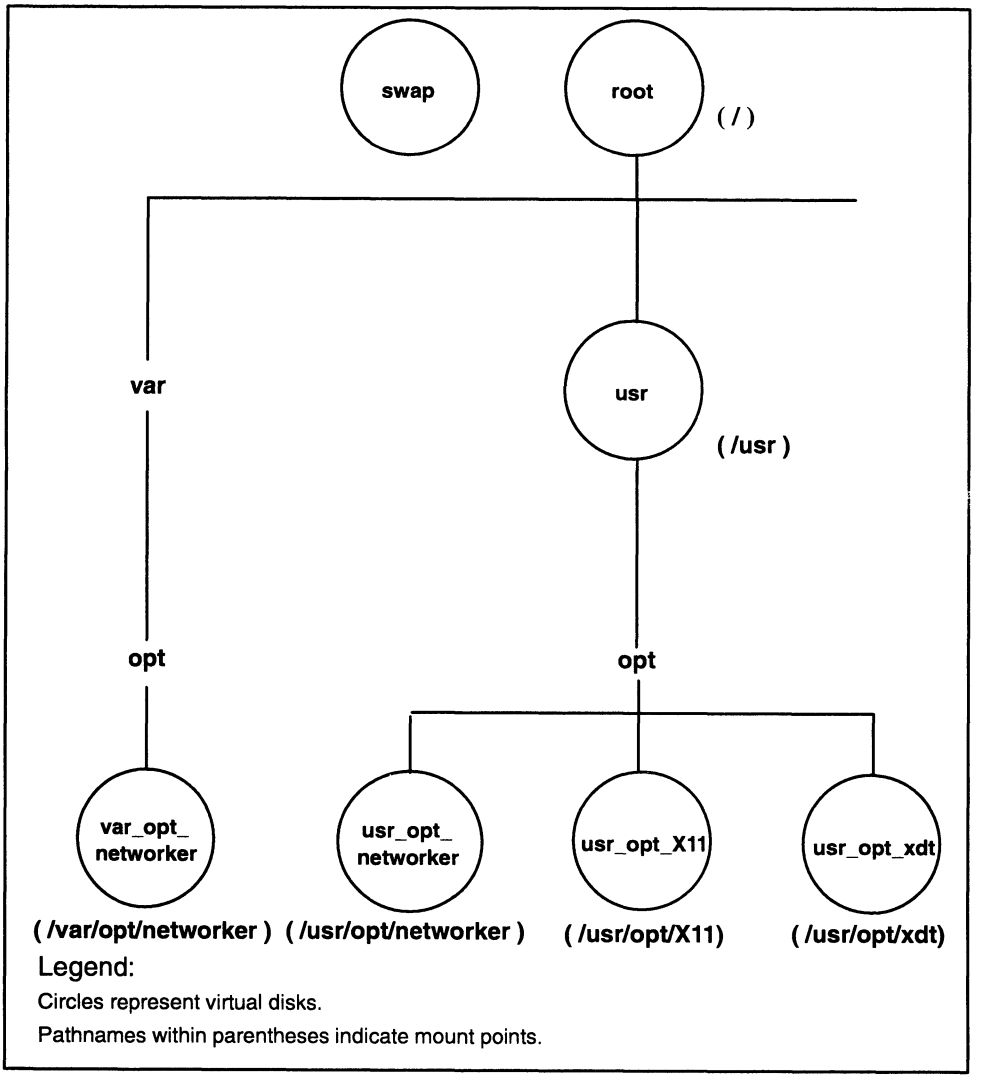

Figure 9–2 User's Perspective of a File System

The operating system prevents you from inadvertently unmounting a file system that someone is using. For example, if a user changed directory to **/usr/opt/X11/catman**, you would not be able to unmount the **/usr/opt/X11** file system. Similarly, if a user were executing a program in **/usr/opt/X11/bin** (regardless of the user's working directory), you would not be able to unmount **/usr/opt/X11** until the executing process terminated.

You further characterize virtual disks with the name of the file system and a mount point location. Information on this procedure is given in this chapter.

File System Tasks

Sysadm provides operations for managing file systems in the File System menu. This chapter addresses these topics:

- Managing local file systems
- Managing remote file systems
- Backing up and restoring file systems
- Retrieving information about files and file systems
- Managing swap space
- File system checking

Managing Local File Systems

Use the **sysadm** `File System -> Local Filesys` menu to create file systems on local disks, manage the file systems, change their size, make them accessible on the local system and on remote systems, and check their internal consistency.

Two Approaches to Creating File Systems

There are two general approaches to creating file systems on the physical medium.

- In most cases, you will use **sysadm** or the **admfilesystem**(1M) command to create file systems on virtual disks, which are located on physical disk devices. If you intend to create virtual disks, the physical disk must be soft formatted first (described in Chapter 7).

  The soft formatting operation builds a virtual disk information table (VDIT) on the disk to manage multiple virtual disks on the physical disk and to provide bad block remapping capabilities. Bad block remapping is the system's automatic way of tracking and avoiding bad blocks on the disk medium.

NOTE: You do not need to set up bad block remapping on highly reliable devices, such as RAID mirrors or nonvolatile RAM disks (NVRDs).

- Alternatively, if you intend to use a physical disk for storage of only one application or database, you do not need to soft format it. Instead, with the **mkfs**(1M) command you can create one file system over the entire physical disk without creating a virtual disk table or one or more virtual disks. This is typically what you do with diskettes and other small media. You trade the benefits of multiple virtual disks and bad block remapping capabilities for the disk space saved by not creating a virtual disk table.

Go to Chapter 15 for information to create file systems on diskettes.

CAUTION: Regardless of how you create a file system, this operation destroys all data on the virtual disk. If a file system already exists on a virtual disk, creating a new file system will destroy the existing one.

File System Features

It may be confusing to know which file system operation to perform. Table 9–3 lists the **sysadm** file system operations (top row) and what they actually accomplish (left column). A checkmark indicates that the **sysadm** operation does perform that action. You must decide on the features you desire for the file systems you create.

Figure 9–3 Sysadm File System Operations and Associated Actions

Actions Associated with sysadm Operations	sysadm Operations				
	Create Virtual Disk	Create DG/UX File System	Add File System	Mount File System	Export File System
Create virtual disk	✓	✓			
Create DG/UX file system	Optional	✓			
Add entry to **/etc/fstab** file		✓	✓		
Create mount point directory		✓	✓		
Mount file system		✓	✓	✓	
Export file system		✓	✓		✓

Choosing the Right Method for Creating a File System

If you prefer to know exactly the file system's location on a given physical disk, create a virtual disk first, and optionally create a DG/UX file system on it using the Create Virtual Disk operation. See Chapter 7 for these instructions. Later, you must perform the Add File System operation to mount the file system and optionally export it using the instructions in this chapter.

If you do not care where the file system is to be located, you may create a DG/UX file system and an implicit virtual disk in a single operation—Create DG/UX File System. This operation includes actions to mount the file system and optionally export it.

You may choose the Add File System operation under these conditions: 1) a DG/UX file system structure has already been created, but you need to change its mount point. Should you physically disconnect a disk from one hardware configuration and transport it to another configuration, you must remount it. 2) Use this operation to identify a non–DG/UX file system for a compact disk (**cdrom**), diskette (**dos**), or a memory file system (**ramdisk**). You may also use this operation for a DG/UX file system; however, a DG/UX file system is implied, by default, in all other file system operations. The file system identifier is included in the **/etc/fstab** file.

The Mount File System and Export File System operations are useful to make them accessible, locally and remotely, following initial creation.

Creating a Virtual Disk and a DG/UX File System in One Operation

Use the **sysadm** operation `File System -> Local Filesys -> Create` to perform all the steps required to create a virtual disk and a file system in one operation.

NOTE: When creating a virtual disk and a DG/UX file system in a single operation, you do not have any control over the virtual disk's physical location. To perform this operation, the system locates and uses whatever free space is available. If you care where your resources come from, do not use this operation. Instead, explicitly create a virtual disk first. See Chapter 7 for instructions.

This operation creates only DG/UX file systems on virtual disks. If you want to add a DG/UX file system that already exists on a virtual disk, or if you want to add a non-DG/UX file system, you should use **sysadm** operation `File System -> Local Filesys -> Add` instead.

You must answer the following queries when you create a local file system:

`Virtual Disk`
: Select a virtual disk from the list, or enter the name of a new virtual disk.

  If you choose to create a new virtual disk, the operation presents the Number of Blocks query.

`Mkfs Options`
: Enter any **mkfs** options you want applied to the new file system.

  See the **mkfs**(1M) manual page for more information about these options.

`Number of Blocks`
: This prompt is displayed only if you create a virtual disk from scratch. Enter the number of blocks for the new virtual disk.

  You cannot choose where the new virtual disk is placed on your physical disk. Use the `Device -> Disk -> Virtual` operations to make any desired changes to the disk.

`Mount Directory (optional)`
: Enter the directory on your system where you want the file system to appear.

  If the directory does not already exist, the Create operation will create it for you.

`Exportable`
: This prompt appears when adding any type file system. If your system is not on a network, this feature does not apply to you.

  Choose whether or not to export the file system, making it available to other systems on your network. Selecting this attribute adds the file system's name to the **/etc/exports** file.

  If you choose to export the file system, the operation presents the Export Options query:

Export Options

Enter any export options for the file system.

You may enter any of the following options, separating them with commas:

secure Require clients to use a more secure protocol when accessing the directory.

NOTE: Secure RPC using DES Authentication is an additional feature that must be purchased separately from the DG/UX ONC/NFS product. You must have this feature to use the **secure** option.

ro Export the directory read-only. If you do not specify the **rw** (read-write) option or the **ro** option, the directory is exported **rw**.

rw=*hostname*[**:***hostname*]...
Export the directory read-only to any systems except those specified in this option. Systems specified in this option have read-write access to the directory. Allowing another system read-write access to the directory does not override normal file and directory permissions. If you do not specify the **rw** option or the **ro** option, the directory is exported read-write to all.

anon=*uid*
If a request comes from an unknown user, use *uid* as the effective user ID. Superusers (UID 0) are considered **unknown** by the ONC/NFS server unless they are included in the **root** option below. The default value for this option is **–2**. Setting the value of **anon** to **–1** disables access by unknown users.

root=*hostname*[**:***hostname*]...
Give superuser access only to superusers from the specified hosts. By default, superusers from other systems do not have superuser access to the directory. A superuser is any user whose UID is 0 (**root** and **sysadm**, for example).

access=*client*[**:***client*]...
Give mount access to each client listed. A client can either be a host name or a net group (see the **netgroup**(4) manual page). Each client in the list is first checked for in the **/etc/netgroup** database and then the **/etc/hosts** database. The default value allows any system to mount the given directory.

NOTE: The create operation will destroy any data currently stored on the virtual disk on which you are creating a file system.

After you answer these queries, the Create operation creates the virtual disk (if needed), the file system, and the mount point directory (if needed). The operation then mounts the file system and optionally exports it.

Adding a Local File System

The system recognizes only the file systems that have entries in **fstab**. When the system boots, it mounts any file systems that have entries in **fstab**. Select the Add operation to put an entry for a file system into the file system table, **/etc/fstab**.

The Create File System operation also adds a file system to the file system table, **/etc/fstab**.

The Add operation's first query is for the file system type. The accepted types are:

`dg/ux` This is the typical disk file system.

`ramdisk` This is a file system that resides entirely in memory, hence the name `ramdisk`, where `ram` stands for random access memory. You do not need to allocate disk space for a `ramdisk` file system. A `ramdisk` file system provides very fast I/O performance for file systems that can fit in memory. The data disappears from memory when you delete or unmount the file system or when the system goes down. See **mfs**(4) for more information.

`cdrom` This is an ISO 9660, High Sierra, or Rock Ridge CD-ROM file system. See Chapter 15 for other CD-ROM considerations.

`dos` This is an MS-DOS file system, created either by **mkfs** with the **dos** flag on an MS-DOS based computer. See Chapter 15 for other MS-DOS file system considerations.

The file system type that you select determines what queries the Add operation presents. The following sections discuss the various queries. The discussion of each query notes the types of file systems to which it applies.

`Virtual disk`

This prompt appears when you add a `dg/ux` type file system.

This is the name of the virtual disk where the file system is located. The virtual disk must already exist. If an existing virtual disk does not appear in the list, it may be because you did not register the physical device.

`Device File`

This prompt appears when you add a `ramdisk`, `dos`, or `cdrom` type file system.

For `ramdisk` type file systems, this is any arbitrary name that you choose to call the device. The name may be any legal file name. When created, the name appears in the **/dev** directory.

For `dos` and `cdrom` type file systems, this is the name of the physical device node file in **/dev/pdsk**.

`Mount Directory`

This prompt appears when adding all file system types.

This is the directory on your system where you want the file system to appear. If the directory does not already exist, the Add operation will ask if it should create the directory.

`Use Wired Memory`

This prompt appears when you add a `ramdisk` type file system.

By default, the system uses unwired memory for the file system. Unwired memory may be swapped out to the swap area on disk in order to free up needed physical memory. If you select wired memory, the file system stays completely in physical memory without being swapped out. Do not select the wired option unless your system has enough physical memory to hold the entire file system and still meet your system's other needs for wired memory. See **mfs**(4) for more information.

`Maximum File Space`

This prompt appears when you add a `ramdisk` type file system.

Set the maximum number of 512-byte blocks that the file system may use. The system does not attempt to allocate memory until the file system requires it. The maximum amount of memory that the system will allocate for the file system is this maximum or the amount of memory available on the system, whichever is less. Any operation that causes the system to attempt to allocate more than this maximum will fail.

`Maximum File Count`

This prompt appears when you add a `ramdisk` type file system.

Set the maximum number of files that the file system may contain.

`Write Permission`

This prompt appears when you add a `dg/ux` or `dos` type file system.

If the write permission is `Read/Write`, users may change the file system as well as read and execute files in it. This is the normal mode for a file system; it does not allow users to override the normal file system security architecture. If the write permission is `Read Only`, users may only read and execute files in the file system.

`Dump Frequency`

This prompt appears when you add a `dg/ux` type file system.

The dump frequency determines how often the file system backup utility, **dump2**(1M), will back up the file system. The backup utility runs whenever you execute **dump2** explicitly or select the **sysadm** operation `File System -> Local Filesys -> Backup -> Create`.

There are four kinds of dump frequency:

`Daily` Back up this file system during every daily, weekly, and monthly backup.

`Weekly` Back up this file system during every weekly and monthly backup.

`Monthly` Back up this file system only during monthly backups.

`None` Do not back up this file system.

You select a dump frequency for a file system based on the importance of the data in the file system and how often you alter it. The schedule that determines when daily, weekly, and monthly backups occur is the backup

cycle, described in "Backing Up and Restoring File Systems" earlier in the chapter.

`Fsck Pass Number`

This prompt appears when you add a `dg/ux` type file system.

The **fsck** pass number, a digit `0` - `9`, indicates the pass on which the file system checker, **fsck**(1M) (when invoked with **fsck –p**), should check this file system for damaged or corrupt files. File systems with pass numbers between (and including) `1` and `9` will be checked in order. That is, all file systems with number `1` are checked first, then those with number `2`, and so on. The digit `0` indicates that a file system should never be checked. This is commonly used for NFS mounted file systems. For more information on **fsck**, see "File System Checking," at the end of this chapter.

`Fsck Logging`

This prompt appears when you add a `dg/ux` type file system.

Select **fsck** logging if you want the file system mounted for fast recovery. With **fsck** logging turned on for a file system, the system logs modifications to the file system. If a failure leaves the file system in a corrupt state, **fsck** can use the log to speed recovery.

Logging is good for file systems where it is crucially important to minimize the amount of time during which the file system is unavailable (as during verification and repair). Because logging has some negative impact on run time write performance in the file system, we recommend it primarily for file systems where rapid recovery and high availability are crucial.

If you select **fsck** logging, the operation later prompts you for log size. Specify the log size in 512-byte blocks. There is a tradeoff in performance between log files of different sizes. A large log file improves run time performance but prolongs recovery time. A small log file degrades run time performance but reduces recovery time. A figure such as 32 blocks or 64 blocks is reasonable.

Configuring your root (/) file system for **fsck** logging is different because you must specify the log size with the ROOTLOGSIZE parameter in the system file and rebuild the kernel. For example, the following line, added to your system file, sets the **fsck** log size for the root file system to 32 blocks:

```
ROOTLOGSIZE    32
```

The next time you boot the new kernel, **fsck** logging will be in effect for the root file system.

`Exportable`

This prompt appears when adding any type file system. If your system is not on a network, this feature does not apply to you.

Choose whether or not to export the file system, making it available to other systems on your network. Selecting this attribute adds the file system's name to the **/etc/exports** file.

If you choose to export the file system, the operation presents the Export Options query.

`Export Options`

This prompt appears when adding any type file system if you elect to make the file system exportable.

See "Creating a File System" for a list of the available options.

After you have answered the queries above, the Add operation presents one more query before executing. If you made the file system exportable, the query is `Mount and export the file system`. If you did not make the file system exportable, the query is `Mount the file system`. Mounting the file system makes it immediately accessible on your system. Exporting the file system makes it immediately accessible to other systems on your network according to any export options you specified.

The Add operation then adds the file system entry to **fstab**. Then, if so requested, it attempts to mount the file system. If the mount directory does not already exist, it asks if it should create it for you before proceeding with the mount. If you chose to export the file system, the operation adds an entry to **/etc/exports** and calls **exportfs**(1M) to make the file system available to other systems in your network.

Deleting a Local File System

Select the Delete operation to remove one or more file systems from the file system table, **/etc/fstab**. If the selected file systems are exportable, the operation also removes their entries from the **/etc/exports** file.

The operation presents a list of file systems and lets you choose which ones to delete. You may also choose whether or not to unmount the file system after deleting it. If the operation cannot unmount the file system (because, for instance, the file system is in use), the operation displays a warning. The operation deletes the file system table entry even if it cannot unmount the file system.

Deleting a file system without unmounting it does not interrupt any work sessions currently using the file system. The next time mounting occurs (at boot or when coming up to run level 2), however, the file system will not be available. The data in the file system remains intact even if no entry for the file system exists in the **fstab** file. To make the file system available again, mount it.

You cannot unmount a file system while users are using it.

You cannot unmount a `ramdisk` type file system if it contains any files or directories. Unmounting a `ramdisk` type file system removes the associated physical device entry in **/dev**.

The operation asks for confirmation before deleting the file system entry.

Expanding a File System

To increase the size of a file system, select the **sysadm** operation `File System -> Local Filesys -> Expand`.

The file system being expanded must be a DG/UX file system that is located on a virtual disk. The operation actually expands the size of the underlying virtual disk, extending the DG/UX file system into the expanded area. The file being expanded may be mounted and online (in use) during this operation.

If you wish to expand a virtual disk that does not contain a file system, go to Chapter 7 for instructions to expand a virtual disk.

To expand the root (/) or **/usr** file systems, remember that you cannot boot from a file system built on a virtual disk that spans multiple physical disks. You need to boot from the root file system because it contains your kernel, **/dgux**. You need to boot from the the **/usr** file system because it contains stand-alone **sysadm**, **/usr/stand/sysadm**, which you may need to boot to recover after a failure. To increase the size of either of these file systems, ensure that the physical disk housing the / or **/usr** file systems contains sufficient space for the creation of another partition. The expansion generally results in an aggregated virtual disk. You may perform this operation while the / and **/usr** file systems are mounted and the system is online and in use.

To expand the size of a mirror, you must unlink its images first. Expand the images separately, then rebuild the mirror. When you rebuild the mirror, you will have to synchronize all but one of the images (the one you use as the master for synchronization).

To expand the size of a cache, you must unlink the cache front–end device(s), and dismantle (uncache) the cache. Expand each virtual disk being cached, and then re–create the cache using the front–end devices that you previously unlinked.

NOTE: You may not use the Expand operation to increase the size of a striped virtual disk.

The operation prompts you for:

`File System to Expand:`
: Specify the mount point directory for the file system to be expanded; for example, **/usr/opt/fredware**.

`Select Space by: [Size alone]`
: You expand a file system by the same method used for creating a virtual disk. Refer to Chapter 7 for a discussion of the methods for creating a virtual disk.

Keep in mind that file system internal data structures and the free space requirement (10% by default) will use some of the space that you are adding to the file system. Consequently, you should add 17% more space than the amount you intend to use. For example, if you need 30,000 more blocks of usable space in a file system, you should add 30,000 + (17% of 30,000) blocks, or 35,100 blocks.

The Expand operation maintains the required percentage of free space for the file system. By default, a file system has a free space requirement of 10%.

Shrinking a File System

To decrease the size of a file system, select the **sysadm** operation `File System -> Local Filesys -> Shrink`.

The file system being shrunk must be a DG/UX file system that is located on a virtual disk on a local host. The operation actually shrinks the size of the

underlying virtual disk and the DG/UX file system. The file system being shrunk must be unmounted and offline (not in use) during this operation.

If you wish to shrink a virtual disk that does not contain a file system, go to Chapter 7 for instructions to shrink a virtual disk.

To shrink the size of a mirror, you must unlink its images first. Shrink the images separately, then rebuild the mirror. When you rebuild the mirror, you will have to synchronize all but one of the images (the one you use as the master for synchronization).

To shrink the size of a cache, you must unlink the cache front–end device(s), and dismantle (uncache) the cache. Shrink each virtual disk being cached, and then re–create the cache using the front–end devices that you previously unlinked.

NOTE: You may not use the Shrink operation to decrease the size of a striped file system.

The operation prompts you for:

`File System to Shrink:`
: Specify the mount point directory for the file system to be shrunk; for example, **/usr/opt/fredware**.

`Number of blocks by which to shrink:`
: Enter the number of 512–byte blocks by which to shrink the file system.

The operation selects which blocks to remove based on internal criteria. The Shrink operation maintains the required percentage of free space for the file system. By default, a file system has a free space requirement of 10%.

Shrinking a file system requires the operation to rearrange data within the file system, collecting unallocated blocks so that the system may free the desired number. The internal structures of a file system complicate this process even more, particularly because some of these internal structures, like the file system itself, are designed to contain a certain amount of free space. The shrink operation will not transgress this built-in cushion of free space because doing so may result in poor I/O performance for the file system. If you specify that the shrink operation recover more space than is available from the file system, it returns an error.

`Percentage of blocks in the remaining DARs to keep unallocated`
: You may set what minimum percentage of the file system should remain empty after the shrinking operation. This figure refers to the percentage of unallocated blocks per DAR (Disk Allocation Region) The default is 20%. If the shrink operation cannot free the desired number of blocks while leaving this percentage of remaining space free, the operation fails.

`Percentage of file node slots in the remaining DARs to keep available`
: You may set what minimum percentage file node slots should remain available per DAR after the shrinking operation. A DAR can hold a fixed number of files and directories (file nodes), and this percentage determines how many of them will remain available for directory or file creation. The

default is 20%. If the shrink operation cannot free the desired number of blocks while leaving the desired percentage of file node slots available, the operation fails.

Under some conditions, the Shrink operation will reduce the file system more than the requested number of blocks. This happens when the Shrink operation causes the new end of the file system to fall in a DAR's metadata area (file node table or free space bitmap). The operation detects this condition and shrinks the file system further, to the beginning of the DAR. This condition does not constitute an error or a problem. When it occurs, the operation prints the following message:

```
The last DAR size was less than the minimal needed to support
the file node table and the DAR bitmap.  The new size of the
file system has been adjusted to nnnn blocks.
```

For a more detailed discussion of the file system internals, see the **mkfs**(1M), **fs**(4), and **tunefs**(1M) manual pages.

Modifying a Local File System

To change the attributes associated with a local file system, select the **sysadm** operation `File System -> Local Filesys -> Modify`. The attributes you can change depend on the type of file system. For more information on the various attributes, see "Adding a Local File System Entry" earlier in the chapter.

You can modify a file system that is in use. After you have specified how you wish to modify the file system, the Modify operation asks if you want to apply the modifications immediately. If you choose to do so, the operation attempts to remount and then unexport or export (as appropriate) the file system so that the changes will be effective immediately.

Displaying Local File Systems

Select the **sysadm** operation `File System -> Local Filesys -> List` operation to display the current file system table (**/etc/fstab**) entries. You may choose to list all entries, only entries for mounted file systems, or only entries for exported file systems.

An example display follows:

```
                                          FS            NFS   Dump  Fsck
File System Source      Mount Directory   Type       RW Mount Frq  Pass
------------------      ---------------   ----       -- ----- ----- ----
/dev/dsk/root           /                 dg/ux      rw        d     0
/dev/dsk/usr            /usr              dg/ux      rw        d     0
/dev/dsk/usr_opt_x11    /usr/opt/X11      dg/ux      rw        d     1
/dev/dsk/usr_opt_aview  /usr/opt/aview    dg/ux      rw        d     1
/dev/ram2               /ram_dir/ram2     dg/ux(ram) rw        x     0
```

The columns in the display correspond to the columns in the **fstab** file. The `File System Source` column shows the physical device pathname. The `Mount Directory` column shows the mount point directory. The `FS Type` column shows the file system type, where the values are:

`dg/ux` A typical disk file system.

`dg/ux(ram)` A `ramdisk` (memory) file system.

`cdrom` The file system on a CD-ROM drive.

`dos` An MS-DOS file system, created by **mkfs** with the **dos** flag.

The `RW` column gives the write permissions mode, either `rw` (read/write) or `ro` (read only). The `NFS Mount` column applies only to remote file systems mounted via ONC/NFS. The field is blank for a local file system. The `Dump Freq` column gives the backup frequency, where possible values are:

`d` Backup during daily, weekly, and monthly backups.

`w` Backup during weekly and monthly backups.

`m` Backup during monthly backups.

`x` Do not back up.

The `Fsck Pass` column shows during which pass the file system checker, **fsck** will check the file system. The value may be `0` through `9`.

For more information on the various **fstab** fields, see "Adding a Local File System" earlier in the chapter.

Mounting and Unmounting a Local File System

If you have unmounted a file system, you have to mount it to make the file system accessible to users. Mount a file system with the **sysadm** operation `File System -> Local Filesys -> Mount`. This operation calls the **admfilesystem**(1M) command, which invokes **mount**(1M).

Remote users to whom you have given access to the file system cannot mount it on their systems until you mount it locally on your own system.

When the Mount operation prompts you for the file system to mount, specify the mount point directory name. If there are no local file systems in your file system table that are currently unmounted, the operation reports this condition before terminating. If an attempt to mount a file system fails, it may be because the file system is corrupted. Verify the integrity of the file system with `File System -> Local Filesys -> Check`.

To make a file system inaccessible, unmount it with the **sysadm** operation `File System -> Local Filesys -> Unmount`. This operation calls the **admfilesystem**(1M) command, which invokes **umount**(1M).

You cannot unmount a file system that is in use. A file system is in use if it contains the working directory for any user or running process, or if it contains a program that is currently running or is open for reading or writing by a program that is currently running.

Exporting and Unexporting a Local File System

To make a local file system accessible to other systems on your network, export it with the **sysadm** operation `File System -> Local Filesys -> Export`.

You can export only file systems that you have made exportable either through the Add operation or the Modify operation. A file system is exportable if the **/etc/exports** file contains an entry for it. The Export operation calls the **exportfs**(1M) command to export the file system.

To make a local file system no longer accessible to other systems on your network, unexport it either with the **exportfs** command (with the **–u** option) or with the **sysadm** operation `File System -> Local Filesys -> Unexport`.

The Export and Unexport operations do not change the **/etc/exports** file, which is essentially a list of file systems for exporting. The Export and Unexport operations simply determine whether or not a file system will be available on remote systems at this time. If a file system appears in the **exports** file, the system makes it available every time the network software starts. The system-maintained file **/etc/xtab** contains a list of currently-exported file systems.

Checking a Local File System

To verify the consistency of a file system, use the **sysadm** operation `File System -> Local Filesys -> Check`.

The Check operation calls the file system checker, **fsck**(1M), to verify the internal structure of the file system. You can also check a file system by performing the Check a File System operation available in **sysadm**'s File System Management Menu.

There are a number of options you can supply to the **fsck** command. The Check operation lets you specify any of these options. For more information on **fsck**, see "File System Checking: fsck," in this chapter. Also see the manual page for **fsck**(1M). For descriptions of **fsck** error messages, see Appendix B.

You can check only unmounted file systems. When the system comes up, it will not mount file systems that are inconsistent and need to be checked. Typically, you need to check file systems only after a failure such as a power outage or disk head crash that has caused the system to halt unexpectedly.

Managing Remote File Systems

The Remote Filesys menu provides operations for managing the remote file system entries in the file system table, **fstab**, and making remote file systems accessible on the local system.

A remote file system appears on your local system to be a file system; for example, you and mount and unmount it. On your system, the remote file system may be any directory, whether or not that directory is the mount point for a file system on its host system. For example, the remote system may have a file system that is mounted as **/usr** and a directory within that file system named **/usr/bin**. You could mount the **/usr** file system from the remote system on your local system as **/usr**, **/usr/bin**, or any other mount directory.

To have access to a remote system's file systems, both the local and remote systems must have the ONC/NFS network package installed and set up. For more

information on ONC/NFS, see *Managing ONC™/NFS® and Its Facilities on the DG/UX™ System.*

The following sections discuss the remote file system operations in more detail.

Adding a Remote File System

If you want users on your local system to have access to a file system from another host on your network, add the remote file system with the **sysadm** operation `File System -> Remote Filesys -> Add`. This operation adds an entry to the file system table, **/etc/fstab**. The Add operation presents several queries, discussed below.

`Mount Directory`
: The directory location on your local system where you want the remote file system to appear. If the directory does not already exist, **sysadm** offers to create the directory.

`Remote Host Name`
: The name of the host where the directory resides. This host name is the name by which the system is known on your network.

  Systems connected to multiple networks have multiple network interfaces; therefore, they have multiple host names. Although you may be able to contact the system using any of the system's host names, it is more efficient (in terms of performance) to use the host name of the network interface connected to your network.

  For example, you want to mount a file system that is located on a system known commonly as **sales03**. This system is connected to two networks and functions as the gateway between the two networks; therefore, the system has two network interfaces. One network interface is called **sales03**, and the other is called **sales03–alt**. Of these two possible host names, you should specify the one that is on the same network as your system. To determine an interface's network, check its Internet address. See *Managing TCP/IP on the DG/UX™ System* for more information on networks.

`Remote Mount Directory`
: This is the directory on the remote system that you want to mount. A file system must be mounted and exported on its native system before other systems in the network can mount it or any directory within it.

`Write Permission`
: If the `Write Permission` is `Read/Write`, local users may change the file system as well as read and execute files in it. You can mount the file system with `Read/Write` permission only if the remote host has exported the file system in **rw** mode. The `Read/Write` mode does not allow users to override the normal file system security architecture. If the `Write Permission` is `Read Only`, local users may only read and execute files in the file system.

`NFS Mount Type`
: If an ONC/NFS file system is hard mounted, user processes will wait indefinitely for file system accesses to complete. This means that if the

remote system on which the file system resides is not responding (because it is down, for example), user programs will appear to hang while they wait for the remote system to become accessible again. You cannot interrupt this kind of hard mount hang unless you specify the `Interruptible` option when mounting the file system.

The benefit of the hard mount option is that it provides more reliability when writing to a remote file system. Hard mounting is preferable for file systems containing valuable data that you update from your remote system.

If an ONC/NFS file system is soft mounted, user processes will time out (terminate with an I/O error) if the remote file system does not respond within a reasonable period of time. The benefit of a soft mount is that the user process does not hang waiting for the remote file system to complete the requested disk I/O. The disadvantage is that there is some risk that a write operation will fail and cause you to lose the data you intended to write. Soft mounting may be preferable for read-only file systems or for file systems containing data that you can afford to lose because of a failed write.

For a more complete discussion of the mount options, see the **mount**(1M) manual page or see *Managing ONC™/NFS® and Its Facilities on the DG/UX™ System*.

`Interruptible`

Set this option to allow you to interrupt remote file system access when it appears hung or becomes inconveniently slow. You interrupt file access the same way you interrupt any other process. For example, the default shell **stty** settings allow you to interrupt a command with **^C**. If you executed the **ls** command in an interruptible, hard-mounted file system whose server had crashed, you would be able to interrupt the **ls** command with **^C**. If the file system were hard-mounted uninterruptible, however, you would not be able to interrupt the **ls** command, and your shell process would hang until the remote server restored access to the file system.

In addition to the interrupt that you can generate using the appropriate control-key sequence from the shell, interrupts may also result when the **kill**(1) command or **kill**(2) system call send a signal to a process or when a modem detects a hangup condition.

The `Interruptible` option is more useful for hard-mounted file systems because accesses to hard-mounted file systems will retry indefinitely until they succeed. In the case of soft-mounted file systems, on the other hand, a failed access will only retry until a finite (and usually brief) time-out period has elapsed.

`Retry in background`

If the remote host does not respond to your system's request to mount the file system, your system retries the mount request a number of times. If attempting a hard mount, your system will continue to retry the mount until the remote host responds. If attempting a soft mount, your system will retry the mount only a few more times.

This query allows you to choose whether the retry occurs in the background or in the foreground. This characteristic is significant when the system is booting because it determines whether the system may continue booting

while a mount attempt retries. By putting mount retries in the background, the system may continue booting (mounting other file systems, starting services, and so on) without waiting on the mount retries.

For more information on how the **mount** command handles failures and retries, see the **mount**(1M) manual page or *Managing ONC™/NFS® and Its Facilities on the DG/UX™ System.*

After you have answered the queries above, the Add operation presents the prompt, `Mount the file system`. If you elect to mount the file system, the operation attempts the mount after adding the file system entry to the **fstab** file. If you select the **soft** ONC/NFS mount type for a read/write file system, the operation presents a confirmation box verifying that you really want the file system soft mounted.

After the Add operation adds the file system entry to the **fstab** file, it attempts to mount the file system if so requested. If the mount directory does not already exist, it asks if it should create it for you before attempting the mount.

Deleting a Remote File System

Select the **sysadm** operation `File System -> Remote Filesys -> Delete` operation to remove one or more remote file systems from the file system table, **/etc/fstab**.

The operation presents a list of file systems and lets you choose which ones to delete. You may also unmount the file systems after deleting them. If an attempt to unmount a file system fails (because, for instance, the file system is in use), the operation displays a warning. The operation deletes the file system table entry even if it cannot unmount the file system.

Removing the file system entry does not interrupt any work sessions currently using the file system. The next time mounting occurs (at boot or when coming up to run level 3), however, the file system will not be available.

Modifying a Remote File System

To change the attributes associated with a remote file system, select the **sysadm** operation `File System -> Remote Filesys -> Modify` operation. For more information on the various attributes, see "Adding a Remote File System Entry" earlier in the chapter.

You can modify a file system that is in use. After you have specified how you wish to modify the file system, the Modify operation asks if you want to apply the modifications immediately. If you choose to do so, the operation attempts to remount the file system if it is mounted.

Displaying Remote File Systems

Select the **sysadm** operation `File System -> Remote Filesys -> List` to display the current file system table (**/etc/fstab**) entries. You may choose to list all entries or only entries for mounted file systems.

An example display follows:

```
                                          FS            NFS   Dump  Fsck
File System Source      Mount Directory   Type       RW Mount Frq  Pass
------------------      ---------------   ----       -- ----- ----- ----
sales03:/pdd/sales03    /pdd/sales03      nfs        rw hard   x     0
sales03:/sales/accounts /sales/accounts   nfs        rw hard   x     0
sales01:/udd/sales01    /udd/sales01      nfs        ro soft   x     0
```

The columns in the display correspond to the columns in the **fstab** file:

`File System Source` The mount directory name on the remote host where the file system resides.

`Mount Directory` The mount point directory on the local system.

`FS Type` The file system type, where the value for a remote file system is always `nfs`.

`RW` The write permissions mode, either `rw` (read/write) or `ro` (read only).

`NFS Mount` How the file system is mounted, `hard` or `soft`.

`Dump Frq` The backup frequency. For a remote file system, this field should always be `x` because you only back up local file systems.

`Fsck Pass` The pass in which the file system checker, **fsck**, will check the file system. For remote file systems, this value should always be `0` because you only check local file systems.

For more information on the various **fstab** fields, see "Adding a Remote File System Entry" earlier in the chapter.

Mounting and Unmounting a Remote File System

If you have unmounted a remote file system, you have to mount it to make the file system accessible to users. You can mount a file system either with the **mount**(1M) command or with the **sysadm** operation `File System -> Remote Filesys -> Mount`.

When the Mount operation prompts you for the file system to mount, specify the mount point directory name, the second field in the **fstab** entry. If there are no remote file systems in your file system table that are currently unmounted, the operation reports this condition before terminating.

If your system's attempt to communicate with the remote system fails, your system will retry the mount request a number of times, possibly causing the Mount operation to appear to hang. For more information on the **mount** command and how it handles request retries, see **mount**(1M). You cannot mount a remote file system unless it is mounted on the remote system.

If you know that the remote system is going to be shut down, or the file system on the remote system is going to be unexported or unmounted, then you should

unmount the remote file systems on your local system. If you do not do this, then some local processes might hang trying to access the remote file system.

To make a file system inaccessible, you unmount it. You can unmount a file system either with the **umount**(1M) command or with the **sysadm** operation `File System -> Remote Filesys -> Unmount`.

You cannot unmount a file system that is in use. A file system is in use if it contains the working directory for any user or running process or if it contains a program that is currently running.

Backing Up and Restoring File Systems

The backup utilities provide a means of saving file systems and restoring files or file systems when needed, such as when recovering after a failure. Typically, you perform a backup operation every evening (or during some off hour), copying to tape any files that have changed since the last backup.

There are several types of backups: monthly (full), weekly, and daily. The default backup cycle defines a one-month schedule for performing backups that involves all three types of backup. First you perform a monthly backup, which copies every file on the system onto the backup tape. Then during the next week, you take daily backups at the end of each work day except Friday. A daily backup copies files changed during the previous day. On Friday, you take a weekly backup, which copies all files changed during the previous week. You continue this pattern for four or five weeks, until the beginning of the next month, when you restart the cycle with another monthly backup.

The reason the backup cycle is such a complicated mixture of backup types is because it makes both file backup and file restoration easier. Consider instead a system where you back up every file, whether it has changed or not, every night. The nightly backup operation would take a lot of time and magnetic media. In addition, restoring a single file would require that you search through the previous night's backup of the entire system. Consider also the opposite scenario, where you make one monthly backup and then supplement it with thirty days of daily backups, for example. The daily backup takes less time, but restoring a file can be quite tedious if you do not remember the exact day when the file was last changed.

In addition to the backup procedures covered in the following sections, you might consider making a bootable tape of your / and **/usr** file systems periodically with the **systemtape** utility. Such a tape can help you recover your system after a failure. See the **systemtape**(1M) manual page for more information.

The Legato NetWorker is also available to manage your backups. See the *Legato NetWorker Administrator's Guide* for more information.

Backing Up Files

You back up systems using the **sysadm** operation `File System -> Backup -> Create` operation. The **sysadm** operation invokes **admbackup**(1M), which invokes the **dump2**(1M) command.

The following sections mention dump levels and dump cycles, terms referring to the schemes used to coordinate dump tapes and the depth of file system changes dumped. For a complete discussion of dump levels and dump cycles, see "Managing the Dump Cycle."

Backing Up with sysadm

To make a backup of a file system, select the **sysadm** operation `File System -> Backup -> Create`. This operation invokes the **admbackup**(1M) command.

The Create operation checks your current position in the backup cycle to see what kind of backup to make: monthly, weekly, or daily. The operation then scans **/etc/fstab** to locate active file systems whose backup-type field matches the backup type scheduled for today according to the backup cycle. The operation then calls the **dump2**(1M) program for each file system that is scheduled to be backed up.

The Create operation prompts you for the following information.

`File System(s)`
: Select the file systems for backing up. Specify **all** to back up all file systems whose backup type (as it appears in **fstab**) matches the backup type for this backup (as indicated by the current position in the backup cycle).

  The **dump2**(1M) command, which performs the backup, does not traverse mount points when backing up a file system. This means that in the process of backing up one file system, it will not also back up another file system mounted within the first. For example, while backing up **/usr**, **dump2** will not also back up **/usr/opt/X11**. If you want to back up **/usr/opt/X11** in addition to **/usr**, you must specify **/usr/opt/X11** explicitly.

  You can back up only local file systems.

`Device`
: Specify the tape or other device to use for the backup. For example, **/dev/rmt/0**.

`Medium Type`
: Specify the type of medium (such as tape) for the backup device. The available types correspond to the various kinds of tapes and other media and the data densities they support. To get a listing of the available media, use the **sysadm** operation `File System -> Backup -> Medium -> List`.

  The List operation gets its information from the readable file **/etc/dumptab**. The default medium type is a QIC-150 cartridge tape. To change the default medium type, use the **sysadm** operation `File System -> Backup -> Medium -> Default`.

`Pack onto One Tape or Medium`
: By default, the system packs a backup tape with as many file systems as possible, to conserve tape. When the tape is full, the system requests that you mount a new tape so it can continue with the backup. As an alternative, this query lets you choose to start each file system backup on a new tape. When the system finishes backing up the file system, it requests that you mount a new tape before starting on the next file system.

`Update Databases`

Select this feature to force the operation to change the backup databases to reflect that this backup occurred. If you select **all** in the `File System(s)` query, the operation updates the backup databases. If you do not select **all** file systems for the backup, the operation leaves the databases unchanged.

`Additional dump2 Options`

Specify any other options for the **dump2** command that the operation uses to perform the backup. You can override the backup level indicated in the backup cycle by specifying a different level as an additional **dump2** option. For example, to specify a full backup of the system, specify the option **–0** (0, zero, represents a full backup). See the **dump2**(1M) manual page for more information.

Backing Up with dump2(1M)

The **dump2** program copies some or all files on a virtual disk to the backup medium based on the backup "level." There are 10 levels: 0 through 9. Execute the **dump2** program by specifying a virtual disk and a backup level as in the following:

```
# /usr/sbin/dump2  -0uf  /dev/rmt/0  /dev/dsk/root ↵
```

where:

`0` Specifies backup level **0**.

`u` Updates the **/etc/dumpdates** file.

`f` Specifies the backup device pathname.

The **dump2**(1M) manual page lists all available options.

The backup level number instructs **dump2** to make a copy of each file that has been modified since the most recent backup at any lower backup level number. For example, if the backup level is 3, **dump2** will make a copy of any file that has been modified since the most recent level 0, 1, or 2 backups. Level 0 backs up every file in the file system because there is no lower backup level. A level 0 backup is often called a *full* backup. A monthly backup is typically a full backup.

The **dump2** command knows that a file has been modified by examining the inode change time (or **ctime**) and the file modification time (or **mtime**) for each file (see **stat**(2) for details). If either of these is later than the backup time for the file system at the appropriate backup levels, then the file has been "modified" since the previous lower level backups. The **dump2** command knows when the file system was last backed up at any given level because it keeps this information in the file **/etc/dumpdates**. This file contains lines of the form:

```
/dev/rdsk/root       0 Fri Jan 14 23:58:58 1994
```

In the example above, the most recent level 0 backup for **/dev/rdsk/root**, the root file system, was made at 11:58 p.m. on January 14, 1994. An entry is added to **/etc/dumpdates** only after the backup completes successfully. This prevents it from inserting a date for a backup that later aborts. Also, duplicate entries are deleted.

In the example above, any other level 0 entries for **/dev/rdsk/root** would be deleted when adding the new one.

Restoring Files

You can restore files using either the **sysadm** operation `File System -> Local Filesys -> Restore` or the **restore**(1M) command. The **sysadm** Restore operation invokes the **admbackup**(1M) command, which invokes the **restore** command; therefore, the two methods are equally reliable.

Restoring Files with sysadm

Use the Restore operation to copy files or file systems from a backup tape to disk. This operation is useful for recovering from disk failures or moving file systems from one disk to another. The Restore operation uses the **restore**(1M) command to retrieve files and file systems from backup tapes created using the **dump2** command.

The following steps outline the simplest, though not necessarily the fastest, method of restoring a file or file system.

1. Before restoring files, make sure the file system where you will restore the files has enough free space to hold them.
2. Retrieve the most recent monthly backup and use it to restore the desired files.
3. Retrieve the weekly backups made since the monthly backup and, loading them in order of earliest first, restore the desired files.
4. Retrieve the daily backups made since the last weekly backup and, loading them in order of earliest first, restore the desired files.

You can shorten the amount of time needed to restore a file if you know when the file was last modified. For example, if on Thursday you accidentally delete a file that you know you modified the previous Wednesday, you need only load Wednesday evening's daily backup in order to restore the file. If you last modified the file last week some time, you only need to load last Friday's weekly backup to restore the file. When you are restoring a group of files or an entire file system, it is more difficult to determine which backups you need, so you may find it easier just to start from the monthly backup and work your way up from there.

Once you have loaded the first tape, you can invoke the Restore operation. The operation prompts you for a directory where it should restore the files. This directory must already exist. The operation also lets you choose between full (noninteractive) mode and interactive mode. Full mode restores an entire file system. Interactive mode is best for restoring individual files because it allows you to search through the tape and restore only the files you want. The interactive mode is managed by the **restore** command. Instructions for using **restore**'s interactive mode are in "Example: Using restore in Interactive Mode," which follows the next section.

Restoring Files with restore(1M)

Restore file systems using the **restore**(1M) command with the **r** option. If a file system (other than / or **/usr**) is completely destroyed, it can be restored by first

remaking the file system and then using the **restore** command on the following tape sets:

1. The most recent monthly backup
2. All weekly backups made since the most recent monthly backup
3. All daily backups made since the most recent weekly backup

NOTE: Attempting to restore the **/usr** file system produces an error when the **/usr/sbin/restore** command attempts to overwrite itself while it is being executed. To solve the problem, copy **/usr/sbin/restore** to **/tmp**, then invoke **/tmp/restore** to restore the **/usr** file system. Alternatively, you may invoke **restore** interactively and select all files in **/usr** except **/usr/sbin/restore** for restoration (see the next section on interactive use). This latter method must be used if you are restoring using **sysadm**.

Consider an example environment where we do our weekly backups on Friday. If the file system is lost on Wednesday of the second week (before the Wednesday backup), we need the following tapes:

Monthly
Weekly (1)
Weekly (2)
Monday (2)
Tuesday (2)

The following steps restore the file system.

1. Unmount the file system:

```
# /etc/umount /dev/dsk/foo ↵
```

2. Remake the file system (be aware that this command destroys all data on the virtual disk):

```
# /usr/sbin/mkfs /dev/dsk/foo ↵
```

3. Check the file system:

```
# /etc/fsck /dev/dsk/foo ↵
```

4. Mount and restore the monthly tape set:

```
# /etc/mount /dev/dsk/foo /mount_name ↵
# cd /mount_name ↵
# /usr/sbin/restore rf /dev/rmt/0 ↵
```

5. Restore the weekly backups and daily backups, one at a time:

```
# /usr/sbin/restore rf /dev/rmt/0 ↵
```

   The **restore r** command restores all files in the current directory.

Example: Using restore in Interactive Mode

You use **restore** in interactive mode either by selecting the Interactive Mode option in the **sysadm** Restore operation or by invoking the **restore** command with the **i** keyletter. In interactive mode, **restore** issues a prompt and waits for you to enter commands. You can use the following commands:

ls	List directory contents (**ls**(1) options are invalid).
cd	Change directory.
pwd	Print working directory.
add	Add file name to the list of files to be extracted.
delete	Delete file name from the list of files to be extracted.
extract	Extract requested files.
quit	Exit program.
help	Print this list.

With the backup tape of the **/sales/accounts** file system mounted, we follow these steps to restore the file **/sales/accounts/smith/redeye**:

1. Change to the directory where **redeye** exists:

```
restore> cd smith ↵
```

2. Verify that **redeye** exists:

```
restore> ls redeye ↵
redeye
```

3. Add **redeye** to the list of files to be extracted:

```
restore> add redeye ↵
```

4. List **redeye** again to verify that it is marked for extraction:

```
restore> ls redeye ↵
*redeye
```

5. Perform the extraction:

```
restore> extract ↵

You have not read any tapes yet.  Unless you know which
volume your files are on, you should start with the last
volume and work towards the first.

Specify next volume #: 1 ↵
Set owner/mode for '.'? [yn] n ↵
```

NOTE: Whenever **restore** asks `Specify next volume #`, always answer "1." This is the correct response, no matter which tape your file is on.

6. Exit **restore**:

```
restore> quit
```

By default, **restore** writes the file to **/tmp** so that you can inspect the file before installing it in its original directory.

After extracting tapes, the operation prompts you to mount the next tape, if desired. You may continue to mount backup tapes and restore files in this manner until you have restored all the files you need.

Managing Backup Media

The **sysadm** operation `File System -> Local Filesys -> Medium` provides operations for managing the backup medium table, **/etc/dumptab**, which contains information about the storage media supported for making backups. The medium table shipped with the DG/UX system probably contains all the media entries you need.

An entry in the medium table includes fields for medium name, the block size used for data transfer, the capacity of each medium (such as a tape cartridge), and a description of the medium. The Add and Modify operations of the Medium menu prompt you for values for each of these fields. The Delete operation prompts you only for the medium name. The List operation lists current medium table entries. An example listing follows:

```
Medium       Block Size  Capacity        Description
------       ----------  --------        -----------
default            16      150M          QIC-150 150MB 1/4" Cartridge Tape
pre_5.4            10      150M          Used for restoring pre-5.4 backups
cartridge_150      16      150M          QIC-150 150MB 1/4" Cartridge Tape
cartridge_320      16      320M          QIC-320 320MB 1/4" Cartridge Tape
cartridge_525      16      525M          QIC-525 525MB 1/4" Cartridge Tape
cartridge          16      150M          QIC-150 150MB 1/4" Cartridge Tape
reel_800           16       19M          800 bpi 1/2" Reel-to-Reel Tape
reel_1600          16       38M          1600 bpi 1/2" Reel-to-Reel Tape
reel_6250          16      143M          6250 bpi 1/2" Reel-to-Reel Tape
reel               16       38M          1600 bpi 1/2" Reel-to-Reel Tape
video              16     2200M          8mm 2GB Video Tape
worm_optimem       16     1200M          2458MB Optimem WORM Platter (1 side)
worm               16     1200M          2458MB Optimem WORM Platter (1 side)
```

The default medium, QIC-150 cartridge, is what appears as the default medium in the Backup menu's Create operation. You may reset this default to any other entry with the Default operation.

Managing the Backup Cycle

Use the **sysadm** operation `File System -> Local Filesys -> Cycle` menu to select a backup cycle, set your position within the selected backup cycle, and list the next backup scheduled for the system.

The backup cycle determines the order and types of backups that occur on your system. The three types of backups are:

Monthly (full) This backup copies every file on the system.

Weekly This backup copies every file changed since the most recent weekly or monthly backup.

Daily This backup copies every file changed since the most recent daily or weekly backup.

The default backup cycle is intended for systems that have a lot of disk space and where a lot of data changes from day to day. The default backup cycle covers a five-week span starting with a monthly (full) backup and followed by a number of daily and weekly backups. After the monthly backup, each of the five following weeks follows the same pattern: there are four daily backups (Monday through Thursday) and one weekly backup (Friday). At the end of the month, you start the cycle over again.

In addition to the default backup cycle, there are two other backup cycles from which to choose. The medium disk cycle performs a complete (full) backup every week, with daily backups the other four working days. This backup cycle is intended for systems with a fairly high amount of disk space but where a relatively low portion of data changes frequently.

The third backup cycle is the small disk cycle. This backup cycle involves a full backup every day. This cycle is good for systems whose file systems can all fit on a single tape.

The default backup cycle, the one intended for large systems, follows this scheme of dump levels:

Dump	Level
Monthly	0
Weekly 1	1
Weekly 2	2
Weekly 3	3
Weekly 4	4
Weekly 5	5
Monday	6
Tuesday	7
Wednesday	8
Thursday	9

The first weekly backup occurs on the Friday following the Friday when you performed the monthly backup. You could perform these dumps in order by executing the **sysadm** operation `File System -> Backup -> Create` every weekday evening, or you could perform them by invoking **dump2** every weekday evening. For example, if you wanted to back up the root file system according to this cycle, dumping it to tape at **/dev/rmt/0**, you would use the following command lines on successive weekday evenings:

Dump	Command
Monthly	**/usr/sbin/dump2 -0uf /dev/rmt/0 /dev/dsk/root**
Monday (1)	**/usr/sbin/dump2 -6uf /dev/rmt/0 /dev/dsk/root**
Tuesday (1)	**/usr/sbin/dump2 -7uf /dev/rmt/0 /dev/dsk/root**
Wednesday (1)	**/usr/sbin/dump2 -8uf /dev/rmt/0 /dev/dsk/root**
Thursday (1)	**/usr/sbin/dump2 -9uf /dev/rmt/0 /dev/dsk/root**
Weekly (1)	**/usr/sbin/dump2 -1uf /dev/rmt/0 /dev/dsk/root**
Monday (2)	**/usr/sbin/dump2 -6uf /dev/rmt/0 /dev/dsk/root**
Tuesday (2)	**/usr/sbin/dump2 -7uf /dev/rmt/0 /dev/dsk/root**
Wednesday (2)	**/usr/sbin/dump2 -8uf /dev/rmt/0 /dev/dsk/root**
Thursday (2)	**/usr/sbin/dump2 -9uf /dev/rmt/0 /dev/dsk/root**
Weekly (2)	**/usr/sbin/dump2 -2uf /dev/rmt/0 /dev/dsk/root**

...and so on.

Week 5 will not always be needed, depending on the month.

The default backup cycle is the file **/etc/sysadm/dumpcycle**. The file looks like this:

```
@[dwm]  0       n       Monthly Set
[d]     6       n       Week 1 - Monday Set
[d]     7       n       Week 1 - Tuesday Set
[d]     8       n       Week 1 - Wednesday Set
[d]     9       n       Week 1 - Thursday Set
[dw]    1       n       Week 1 - Weekly Set
[d]     6       n       Week 2 - Monday Set
[d]     7       n       Week 2 - Tuesday Set
[d]     8       n       Week 2 - Wednesday Set
[d]     9       n       Week 2 - Thursday Set
[dw]    2       n       Week 2 - Weekly Set
[d]     6       n       Week 3 - Monday Set
[d]     7       n       Week 3 - Tuesday Set
[d]     8       n       Week 3 - Wednesday Set
[d]     9       n       Week 3 - Thursday Set
[dw]    3       n       Week 3 - Weekly Set
[d]     6       n       Week 4 - Monday Set
[d]     7       n       Week 4 - Tuesday Set
[d]     8       n       Week 4 - Wednesday Set
[d]     9       n       Week 4 - Thursday Set
[dw]    4       n       Week 4 - Weekly Set
[d]     6       n       Week 5 - Monday Set
[d]     7       n       Week 5 - Tuesday Set
[d]     8       n       Week 5 - Wednesday Set
[d]     9       n       Week 5 - Thursday Set
[dw]    5       n       Week 5 - Weekly Set
```

The columns in the table are:

`Cycle` The first column lists the cycle letters that correspond to those in **/etc/fstab**, which indicate when the file systems will be backed up. For instance, file system **/comm** might be set to `w`, so it is only backed up once per week. The cycle letters are:

- `dwm` Archive daily, weekly, and monthly.
- `wm` Archive weekly and monthly.
- `d` Archive daily only.
- `w` Archive weekly only.
- `m` Archive monthly only.
- `x` Do not archive at all.

`Level` The second column shows numbers that are used internally by the **dump2** program. The ones we supply need not be changed for normal system operation.

`Multi` The third column indicates whether multi-dumping shall be in effect for the backup. Multi-dumping means backing up more than one file system per tape. If `y`, multi-dumping occurs. This means as many file systems as there is room for will be written to tape. An `n` entry means write just one file system per tape.

`Description`
This column is a comment describing the backup cycle entry. We recommend that you label your tapes so that they match the entries in the backup cycle list. `Monthly` means backup all file systems marked with `d`, `w`, or `m` (daily, weekly, or monthly). `Monday Set` means backup all file systems marked with `d` (daily), and so on with the other weekdays. Friday's backup becomes part of the `Weekly Set` of backup tapes.

The "at" symbol (@) indicates the current position in the backup cycle.

The complete set of backup cycles that provided with the DG/UX system are in the directory **/etc/sysadm/dumpcycles**. This directory also contains descriptions of each backup cycle.

To set a backup cycle for use on your system, select the **sysadm** operation `File System -> Backup -> Cycle -> Select`.

The system keeps track of your current position in the backup cycle. After every backup, the system moves the pointer (indicated by @) ahead to the next backup, advancing line-by-line down the cycle until the end of the month. On the first day of the next month, use the Position operation to reset the pointer to the top of the list to restart the backup cycle.

It is possible for the current pointer position in the list to be wrong if backups are skipped for a day or more. To restore the backup cycle pointer to the correct position, use the Position operation.

To display your current location in the backup cycle, select the List operation.

For a complete description of backup cycle format, see **dumpcycle**(4).

Making Tapes

This section does not discuss a **sysadm** menu procedure; it simply offers some suggestions for making tapes. When you have a small-scale backup task, as when you are making a personal tape for a user, you don't need to use the **dump2** and **restore** operations. You can use **cpio**(1) instead. See the **cpio** manual page for a complete listing of options and further instructions. To backup a directory named **/sales/smith** (and all subdirectories and files under it), do the following:

1. Mount a tape and put the drive on-line.
2. Go to the directory you wish to backup:

   ```
   # cd /sales/smith ↵
   ```

3. Backup everything in the directory to tape:

   ```
   # find . -print | cpio -ocvB > /dev/rmt/0n ↵
   ```

   where *n* =

   n No rewind

   l Low density

   u Uncompressed

   The contents of **/sales/smith** have been backed up to tape. The **cpio** options we used are:

   o Copy files to standard output.

   c Use ASCII headers for portability.

   v Be verbose: print a list of file names.

   B Use large block size: 5120 bytes instead of 512.

   To write individual files to tape, go to the directory where the files are located and type:

   ```
   # echo fileA fileR fileZ | cpio -ocvB > /dev/rmt/0 ↵
   ```

   This command backs up the contents of all three files.

   We directed the output of the backup to raw magnetic tape (**rmt**), device 0.

Retrieving Information about Files and File Systems

This section shows you how to find and display information about files in your file systems. The **sysadm** `File System -> File Information` menu has three operations for getting information about files and file systems:

`Disk Use` This operation displays information about disk space taken up by file systems.

`Check` This operation searches for files that may constitute a security risk: device files not located in **/dev** and executables owned by the superuser that have the setuid bit set.

`Find` This operation locates files and directories based on criteria that you specify.

We use the following terms in this section:

inode Data structure containing information about a file such as file type, size, date of creation, owner ID, and group ID. The number of inodes represents the total number of files that can exist on the system. The **mkfs**(1M) program, which creates a file system, accepts options that you can use to control the number of inodes in a file system. An inode is 126 bytes long, and there are 4 inodes to a disk block.

disk block A 512-byte unit of data as it is actually stored and manipulated.

setuid A mode bit that can be specified for any executable file. When a user runs an executable file that has the setuid bit set, the system gives the user the permissions of the owner of the executable file for the duration of the command. See **chmod**(1).

/dev Administrative directory containing entries for all devices on the system.

Displaying Disk Space Usage

Select the **sysadm** operation `File System -> File Information -> Disk Use` to display a table showing the number of blocks and inodes in use on mounted file systems that you specify. You may enter multiple file system names. If you specify no file system names, the operation lists information for all file systems. Here is an example of a display that Disk Use provides:

```
            Free      Total     Pct      Free      Total     Pct
Directory   Inodes    Inodes    Used    Blocks    Blocks    Used
---------  -------  --------  ----  --------  --------  ----
/tmp/root     5147      5760   10%     24618     40000   38%
```

Checking for Security Breaches

Select the **sysadm** operation `File System -> File Information -> Check` to search a directory tree for files that have suspicious ownership and permission settings. If you specify no directories, the operation checks the system's entire directory tree. This operation may be time consuming.

The Check operation finds files that may indicate that a security breach has occurred. This command searches the directory you specify and reports files that have the following problems:

- Device files that exist outside of **/dev**. No device files should reside outside **/dev** unless you, as system administrator, have created or moved device files for a special purpose, such as a test.
- Nonsystem files that are owned by the superuser (user with UID of 0, **sysadm** and **root** by default) and have the setuid bit set.

The setuid bit is a special kind of permission attribute that all files have but which is only useful for executable files (such as programs and shell scripts—not directories or text files). When a user runs an executable that has the setuid bit set, the process runs with the effective user ID of the executable's owner—not, as is normally the case, with the user ID of the person running the executable.

For example, if user **fred** executes a program owned by user **bob**, and this program does not have the setuid bit set, **fred**'s process runs with the effective user ID of **fred** (as normal). On the other hand, if user **fred** executes a program owned by user **bob**, and this program does have the setuid bit set, **fred**'s process runs with the effective user ID of **bob**. This means that the program, if it is so written, can access files to which **fred** may not normally have access, but to which **bob** does. User **fred** does not even need to know **bob**'s password for these accesses to occur.

The rationale behind the setuid bit is to allow users to perform some action that they should not be able to perform under normal circumstances. The **lp** command, for example, has the setuid bit set so that any user who invokes **lp** to queue up a print job can, through the agency of the **lp** program, do things such as copy files to the LP system's directories and add requests to the LP scheduler's queue file.

When you execute a program that has the setuid bit set, your actions are determined by the scope of the program and the user ID of the program's owner. Thus, we arrive at the danger inherent in the setuid bit feature: the combination of a permissive program that has the setuid bit set, owned by a privileged user ID, may give a user too much freedom on the system. This situation can result in a breach of security. The extreme case would be a shell program owned by **root** that has its setuid bit set; any user could execute the program and enjoy the use of a superuser shell throughout the system.

Among its various functions, the Check operation includes a search for suspicious programs, ones owned by **root** that have the setuid bit set.

The following example file listing shows **ls –l** output for several files that have the setuid bit set (which we know because of the **s** flag in the group-execute permission and/or owner-execute permission places in the permissions line). The file **/sales/tom/nasty** is suspicious because it is owned by **root** but is obviously not a normal system program. It appears to belong to user **tom**. Depending on what Tom's program does, it may constitute a security breach.

```
-rwsr-sr-x  1 root  bin    44924 Mar 28 18:16 /usr/bin/at
-rwsr-sr-x  1 root  bin    29500 Mar 28 18:16 /usr/bin/crontab
---s--x---  1 root  users  95376 Aug 18 11:08 /sales/tom/nasty
```

Tom would never have been able to create such a program without superuser access. Thus, the program may not only constitute a security breach itself, but it also indicates the presence of a breach elsewhere, the one that allowed Tom to become superuser and produce the suspicious executable.

To protect your system, never leave your logged-in terminal unattended (particularly if you are logged in as **root** or **sysadm**). Another user could move or copy files, or commit any manner of destructive acts, all with your user ID.

When you find a setuid bit set, investigate further. You may need to correct setuid permissions with the **chmod**(1) command. In general, you will not be creating device files, so none should exist outside of **/dev**. There might, however, be the case when you or someone else creates a test device file outside of **/dev**. If necessary, you may either move or delete the device file.

For environments where you require a greater degree of security, there are the Trusted DG/UX systems, which provide B1 and C2 levels of security. For more information on the Trusted DG/UX systems, contact your Data General representative.

Finding Files

Select the **sysadm** operation `File System -> File Information -> Find` operation to search a specified directory and list all files or directories under it that satisfy specified criteria. The operation can also sort the output data in various ways. (In this discussion, the term *file* also refers to directories.)

The Find operation is useful as a part of everyday maintenance. You can use it to find

- files whose names match a specific pattern,
- files belonging to particular users or groups,
- files of a given type,
- files that have not been accessed or modified in a long time and are no longer necessary, or
- files that take up too much space.

You can also determine the number and order of files found. You can sort the information by name, size, access date, and modification date.

The Find operation presents you with the following queries.

`Directories`
: List the directories you want to search. The operation searches the named directories as well as any directories underneath them. If you specify no directories, the operation searches the entire system.

`Restrict to Local File Systems`
: Select this option to restrict the search to file systems that reside physically on your system. Without this option, the operation searches local directories as well as file systems mounted from remote systems.

`Restrict to This File System`
: Select this option to restrict the search to the file systems containing the named directories. Without this option, the operation searches all directories under the named directories as well as any file systems mounted under the directories.

File Name

Specify name or name pattern to match. You may use the metacharacters (wildcards) accepted by **sh**(1). These metacharacters are:

? Matches any one character.

***** Matches any number of any characters.

[ ] When surrounding a group of characters (not including the hyphen,–), matches any one character in the group. For example, **[abc]** matches an occurrence of **a**, **b**, or **c**. Use the hyphen to represent a range of characters. For example, **a–z** represents any lowercase letter. Follow the closing bracket with **!** to negate the set, causing the expression to match any character *not* in the set. For example, **[ag–iA12]!** matches any character except **a**, **g**, **h**, **i**, **A**, **1**, or **2**.

If the name contains metacharacters, surround it with quotation marks. For example, the pattern "**c*.[1–4ab]**" matches any file or directory name starting with **c** and ending with a dot followed by one of the characters **1**, **2**, **3**, **4**, **a**, or **b.**

Owner Name or ID

Specify the login name or user ID number of an existing user. The operation finds only files that the user owns. Remember that usernames are case sensitive.

Group Name or ID

Specify the group name or group ID number of an existing group. The operation finds only files whose group ownership is for the specified group. Like usernames, group names are case sensitive.

File Type

Specify the type of file to find. You may specify more than one type. The DG/UX file system and the Find operation recognize these file types:

`any`	Any of the following types.
`regular`	A typical file such as a text file that is not a directory or any other type specified here.
`directory`	A directory.
`block special`	A device file created for block data access.
`character special`	A device file created for character (raw) data access.
`fifo (named pipe)`	A named pipe.

Days Since Last Modification

When a user modifies a file, the system records the date in the file's inode, the block containing data about the file. The Find operation can use this information to search for files that have not been modified since a particular date.

Days Since Last Access

A file's inode also includes the date the file was last accessed (read). The Find operation can use this information to search for files that have not been accessed since a particular date.

`File Size (bytes)`
: Specify a size limit for files. The operation searches for files larger than this size.

`File Sorting Method`
: The operation lets you organize the data about files that it finds. You may choose to sort files by name, size, access time, or modification time. You can also specify increasing order or decreasing order for the sort, or you can specify no sorting at all.

`Maximum Number of Files to Report`
: To limit the number of files in the display, specify an upper limit. Specifying zero removes the upper limit.

Managing the Swap Area

A swap area is not a file system, but the DG/UX system manages swap areas in much the same manner as file systems. For each swap area on your system, there is an entry in your file system table, **/etc/fstab**. Depending on how you configured your system, the system may check **fstab** at boot for any swap entries and use the listed virtual disks as swap areas. The Swap Area menu provides operations for managing swap area entries in **fstab**.

If you wish to use two physical disk areas for swapping, you should create a second virtual disk and add it as a swap area rather than creating a two piece virtual disk and swap on that one virtual disk. This is because the swapping code in the kernel will take advantage of the two areas to optimize swapping operations. If the two areas are bundled as a single virtual disk, the swapping code does not recognize they are separate and might not spread the swapping load between them.

Adding Swap Area

By default, your system has one swap area 24 Mbytes (50,000 512-byte blocks) in size. Your swap area may be different if your system is an OS client or if you changed the default swap size during installation. If you know that your applications will need additional swap area, you may add more.

There is no universally applicable formula for calculating how much swap area you need. You need to consider how much memory your applications require and how much physical memory your computer has. A good general guideline is to start with swap area equal to 1.5 times your physical memory. For example, a system with 16 Mbytes of memory should have 24 Mbytes (50,000 blocks) of swap area.

You know you have too little swap area when processes terminate unexpectedly and messages like this appear on your system console:

```
From system:  Out of paging area space
```

A system may have swap space on its own disks, or it may have swap space on another system accessible over the network. You may not have swap space on a local disk as well as on a remote host's disk. Typically, a system that has its own disks with its own OS software will also have its swap area on a local disk, while an OS

client will have its swap area on the server's disk. If any OS client has its own disk, it may create its swap area there even though its OS software comes from an OS server on the network. For more information on the local swap/remote OS configuration, see *Customizing the DG/UX™ System.*

NOTE: You are discouraged from using a nonvolatile random access memory (NVRAM) device for swap space. A NVRAM device does not have enough memory to be a reasonable swap resource.

To add local swap space in addition to the default swap area already established for your system, use the **sysadm** operation `Device -> Disk -> Virtual -> Create` to create a virtual disk of the desired size. You should not create a file system on a virtual disk to be used for swap space.

After you have created the virtual disk, add the swap area to **fstab** and make it usable on the system by selecting the **sysadm** operation `File System -> Swap Area -> Add`.

The Add operation lists the available virtual disks, the ones not already appearing in **fstab**, and lets you choose one. It then adds the **fstab** entry and calls **swapon**(1M) to make the new swap area active immediately.

You cannot make a swap area inactive while the system is running. Once you have made a swap area active, it remains active until you shut down the system.

To make a swap area active every time you boot the system, you not only need to add an **fstab** entry, but you also need to change your system parameters so that the **swapon**(1M) program, when executed by the system at boot time, will run with the **-a** argument. Make this change by selecting the **sysadm** operation `System -> Parameters -> Set`. In the query for `arguments to swapon`, specify **-a**. The **-a** option to **swapon** causes it to mount all swap areas listed in the **fstab** file at boot time.

Deleting Swap Area

To remove a swap area entry from the **fstab** file, select the **sysadm** operation `File System -> Swap Area -> Delete`. Deleting a swap area entry does not immediately stop the system from using the swap area. The system will continue to use the swap area until you reboot.

Displaying Swap Areas

To list the swap areas in your file system table, select the **sysadm** operation `File System -> Swap Area -> List`. A sample swap area display appears below:

```
Swap Areas
--------------------------
/dev/dsk/swap
/dev/dsk/swap2
```

The display shows the pathnames of virtual disks entered in the file system table as swap area.

File System Checking

The **fsck** program checks and repairs file systems (see the **fsck**(1) manual page for complete details). The **fsck** program can check file systems in either of two modes, depending on the kind of file system:

`Multi-pass fsck`
: Checks normal file systems by performing five passes through the file system, each pass verifying different aspects of the file system.

`Fast recovery fsck`
: Performs single-pass verification of file systems that were mounted to take advantage of the **fsck** logging feature, described in the next section.

You may invoke **fsck** two different ways. Select the **sysadm** operation `File System -> Local Filesys -> Check` or invoke **fsck** directly at the shell prompt. For more information, see "Invoking the fsck Program" later in the chapter.

You need to check a file system only after a system crash, disk or disk controller hardware failure, or any other time when an abnormal event may have caused service to the file system to terminate unexpectedly. Normally, the system checks file systems as necessary when it boots.

You must use this program when you are bringing up your system after an abnormal shutdown such as a power outage or system crash.

The **fsck** program checks blocks and file sizes, directory contents, connectivity, link counts and resource allocation, and disk allocation region (DAR) information, including the free-block bitmap, the free-inode list, and summary counts. Multi-pass **fsck** checks these features in this order. Fast recovery **fsck** does not necessarily check them in this order. Both modes of **fsck** report inconsistencies that they find. It is your option to fix the inconsistencies or ignore them.

For more information on the errors and messages that **fsck** returns, see Appendix B.

This section:

- Discusses the normal updating of the file system.
- Discusses the possible causes of file system corruption.
- Presents the corrective actions taken by **fsck**. It describes both the program and the interaction between the program and the system administrator.

Fast Recovery File Systems

You can reduce the amount of time that **fsck** requires to check a file system by requesting **fsck** logging for the file system. For fast recovery file systems, the system keeps a log called an intent log containing records about file system updates. If a failure occurs and you need to restore the file system, **fsck** can use the log to speed the process of verifying and restoring file system integrity.

Logging is good for file systems where it is crucially important to minimize the amount of time during which the file system is unavailable (as during verification

and repair). Because logging has some negative impact on run time write performance in the file system, we recommend it primarily for file systems where rapid recovery and high availability are crucial.

Any file system can be a fast recovery file system. You specify **fsck** logging for a file system when you mount it. Mount a file system with the **sysadm** operation `File System -> Local Filesys -> Mount`.

Select the **Fsck Logging** feature when prompted. The operation will later prompt you for the log size. If you edit the **/etc/fstab** file yourself, you add the **fsck_log_size** option in the options field, specifying a size for the **fsck** log as in the following sample **fstab** entry:

```
/dev/dsk/sales /sales dg/ux rw,fsck_log_size=64 d 1
```

You specify the log size in 512-byte blocks. There is a tradeoff in performance between log files of different sizes. A large log file improves run time performance but prolongs recovery time. A small log file reduces recovery time but degrades run time performance. Recovery time depends on the size of the intent log, not the size of the file system.

To configure your root (/) file system for **fsck** logging, specify the log size with the ROOTLOGSIZE parameter in the system file and rebuild the kernel. For example, the following line, added to your system file, sets the **fsck** log size for the root file system to 32 blocks:

```
ROOTLOGSIZE    32
```

The next time you boot the new kernel, **fsck** logging will be in effect for the root file system.

File System Update

Every time a file is modified, the DG/UX operating system performs a series of file system updates. When written to disk, these updates yield a consistent file system.

There are five types of file system updates. The updates involve the following areas:

- the superblock
- inodes
- index (indirect) blocks
- data blocks (directories and other files)
- disk allocation region information, which includes the free-block bitmap and the inode table

Corrupted File Systems

Many things can corrupt a file system. Improper shutdown procedures and hardware failures are the most common causes.

Some examples of improper shutdown procedures are:

- Forgetting to use the **shutdown**(1M) command (which unmounts all file systems, including the root) before halting the CPU.
- Physically write protecting a mounted file system.
- Taking a mounted file system off line.

Each DG/UX file system contains a flag in the superblock which indicates whether or not the file system is mountable. You can only mount a file system if it is marked as mountable; if a file system is unmountable, you must first run **fsck** in order to repair inconsistencies. The **fsck** program will mark a file system mountable only when it is internally consistent. If an attempt to mount a file system fails with the message `No space left on device`, the file system is probably corrupt and needs to be checked.

A file system is marked mountable when it is created. It is marked unmountable whenever it is mounted, and is not marked mountable again until it is either unmounted or has passed the **fsck** program's internal consistency checks. Therefore, file systems which were still mounted at the time of an abnormal system shutdown cannot be remounted until **fsck** has been run over them, whereas those file systems which were cleanly unmounted before shutdown can immediately be remounted.

Fixing Corrupted Files

This section discusses ways to discover and fix inconsistencies for different kinds of update requests.

The **fsck** program lets you check a file system for structural integrity by performing consistency checks on redundant data. Redundant data is either read from the file system or computed from other known values. When **fsck** reports an inconsistency, it asks whether the inconsistency is to be corrected by repairing or deleting the corrupted item. You can accept or reject this request. In the following example, **fsck** finds an incorrect link count in Phase 4 and asks if it should fix the problem. The file system being checked was not mounted for fast recovery.

```
# fsck /dev/dsk/mydisk ↵

** /dev/dsk/mydisk:
** Phase 1 - Check Blocks and File Sizes
** Phase 2 - Check Directory Contents
** Phase 3 - Check Connectivity
** Phase 4 - Check Link Counts and Resource Accounting

Inode 67 (owner: 2 [bin]; group: 2 [bin]; size: 52736 bytes;
    type: Ordinary; mode: 755; mtime: Fri Nov 20 17:54:36 1987)
    has incorrect link count (2 should be 1) -- fix?  y ↵

** Phase 5 - Check Disk Allocation Region Information
File system is now mountable.

13936 of 50000 blocks used (36064 free); 288 of 5822 inodes
    used (5534 free).

#
```

If you respond **y**, the **fsck** program corrects the incorrect link count that it found for inode 67 and moves on to Phase 5. When it finishes, it reports that the file system is mountable.

If you respond **n**, **fsck** does not fix the inconsistency, and the file system remains unmountable.

Superblock and Disk Allocation Region Information

The superblock and disk allocation region information are some of the most commonly corrupted items. Every change to the file system's blocks or inodes modifies the superblock and the disk allocation region information. The superblock and disk allocation region are most often corrupted when the system was not properly shut down with the **shutdown**(1M) command.

Superblock and disk allocation region inconsistencies can involve file system size, the number of available inodes and blocks, the free-block bitmaps and the free inode lists. The following sections give brief discussions of these sections.

Free-Block Bitmap

Each disk allocation region (DAR) contains a bitmap representing all the blocks in the DAR. Multi-pass **fsck** compares that information with its own map of allocated blocks, looking for inconsistencies that suggest DAR corruption.

Free-Inode List

Each DAR contains a link list of free inodes in that DAR. The **fsck** program ensures that all free inodes in the DAR appear in the list, and that no allocated inodes appear in the list.

Summary Counts

The superblock and each DAR contain several counts: the number of used inodes, the number of used blocks, the number of directories. The **fsck** program compares these counts to the information it has compiled.

Inodes

An individual inode is less likely than the superblock to be corrupted. However, because of the great number of active inodes, the free inode lists are as susceptible as the superblock to corruption. Multi-pass **fsck** checks for inconsistencies involving format and type, link count, duplicate blocks, and inode size. Fast recovery **fsck** checks for inconsistencies involving link count, duplicate data element pointers, and inode size.

Format and Type

Each inode contains mode information. This information describes the type of the inode. Inodes may be one of eight types: regular, directory, control point directory, special block, symbolic link, special character, FIFO, or socket. Any other type is illegal.

Link Count

Each inode contains a count of the directory entries linked to the inode. Multi-pass **fsck** verifies the inode count by checking down the total directory structure, starting

from the root directory, and calculating an actual link count for each inode. Fast recovery **fsck** verifies link counts by comparing them to link count records in the log.

In multi-pass checking, when the link count (which is stored on the disk) is nonzero and the actual link count (kept by **fsck**) is zero, no directory entry appears for the inode. If no entry appears, **fsck** may link the disconnected file to the **lost+found** directory. Fast recovery has records that enable it to match lost files with their original directories; consequently, it does not leave files or file fragments in **lost+found**.

If the stored and actual link counts are nonzero and unequal, **fsck** may replace the link count on the disk by the actual link count. When this situation arises, a directory entry may have been added or removed without the inode being updated.

Duplicate Blocks

Each inode contains a list and sometimes pointers to lists (index blocks) of all the blocks claimed by the inode.

Multi-pass **fsck** checks these lists for duplicate blocks. Duplicate blocks can occur when a file system uses blocks claimed by both the free-block bitmap and other parts of the system or when two or more inodes claim the same block. Any block claimed more than once is flagged by **fsck** as a duplicate block. If there are any duplicate blocks, **fsck** makes a partial second pass of the inode list to find the inode of the duplicated block. If the files associated with these inodes are not examined for correct content, **fsck** will not have enough information to decide which inode is corrupted and should be cleared. Usually, the inode with the earliest modification time is incorrect and should be cleared.

Fast recovery **fsck** does not check for duplicate blocks. It compares inodes and index blocks with log records listing which blocks were allocated for each file.

Size Checks

Each inode contains a size field. This field's size indicates the number of bytes in the file associated with the inode.

The **fsck** program can check the size for inconsistencies, such as directory sizes that are not a multiple of 512 bytes, or a mismatch between the number of blocks actually used and the number indicated by the inode size. The **fsck** program also checks for directory corruption, where conflicting information is found within the directory entries.

The **fsck** program can also perform a check of the size field of an inode. Multi-pass **fsck** uses the size field to compute the number of blocks that should be associated with the inode, and then compares that number to the actual number of blocks claimed by the inode. Fast recovery uses log records to track changes to the file size and number of blocks claimed by the inode.

Control Point Directories

The root inode of a file system is a special type of directory known as a *control point directory.* A control point directory is like an ordinary directory except that it has resource limits associated with it for inodes and for data blocks. Effectively, these

limits determine how many files you may create in the file system and how large the file system (total size of all files and directories) may become. The total resources consumed by the control point directory and all its descendants (to which it is the *space parent*) may not exceed the size limits. The limit on data blocks (size) is determined by the size of the virtual disk on which the file system resides. To change the size of the file system, use the Expand and Shrink operations in the **sysadm** `File System -> Local Filesys` menu.

Index Blocks

Index blocks (also known as indirect blocks) are owned by an inode. Therefore, inconsistencies in an index block directly affect the inode that owns the block.

Multi-pass **fsck** can check inconsistencies involving blocks already claimed by another inode and block numbers outside the range of the file system.

Fast recovery **fsck** uses the log to track changes to the contents of index blocks.

Data Blocks

There are two types of data blocks: plain data blocks and directory data blocks. Plain data blocks contain the information stored in a file. Directory data blocks contain directory entries. The **fsck** program does not try to check a plain data block.

There are several checks that **fsck** makes for directory data blocks. Multi-pass **fsck** checks all entries in all directory data blocks, but fast recovery **fsck** checks only those entries that have been modified. The directory data block check searches for:

- bad self-identification information
- directory entries for unallocated inodes
- directory entries for inodes which do not exist in the file system
- directories that are disconnected from the file system

If a directory entry inode number points to an unallocated inode, **fsck** may remove that directory entry. Directory entry inode numbers can point to unallocated inodes when the data blocks containing the directory entries are modified and written out, but the inode is not written out.

If a directory entry inode number is pointing to a nonexistent inode, **fsck** may remove that directory entry. This condition occurs if bad data is written into a directory data block.

The **fsck** program checks that all directories are linked into the file system; i.e., they have a parent directory pointing to them (except for the root). Multi-pass **fsck** links unlinked directories to the file system's **lost+found** directory. Fast recovery **fsck** links unlinked directories back to their original places in the directory tree. When inodes are being written to the file system without the corresponding directory data blocks being written to the file system, the directories are not linked into the file system.

Invoking fsck

You can invoke **fsck** any of five ways:

sysadm Select the **sysadm** operation `File System -> Local Filesys -> Check`.

rc script You can set up any of your system's **rc** scripts to invoke **fsck** when you change run levels. For more information on **rc** scripts and how to set them up, see Chapter 3.

Command line From the command line, type:

fsck [*options*] [*file_system_names*]

where *options* are single character flags that modify the behavior of the command and *file_system_names* refer to the file systems that you want to check. The *options* are covered in detail later in the next section.

Initialization The initialization version of **fsck** is built into the operating system kernel and is automatically run over the root file system when you boot your system.

Stand-alone You boot stand-alone **sysadm**, and select the operation `File System  Check a File System`.

Options to fsck

All options are represented by single-character flags; options must begin with a hyphen. All options except for **–t** are Boolean flags, and may thus be combined. For example, you can combine the options **–p**, **–x**, and **–D** as follows: **fsck –pxD**.

The following options are interpreted by **fsck**:

–l Perform fast recovery using the **fsck** log, if possible. The **fsck** program can perform fast recovery only if:

- You specified the **fsck_log_size** option, which turns on fast recovery logging, the last time you mounted the file system (unless it was the root file system, for which fast recovery logging is in effect if you set the ROOTLOGSIZE parameter when you built the kernel), *and*
- This is the first time that **fsck** has checked the file system since that mount.

If fast recovery **fsck** cannot repair the file system, or if you specify the **–l** option for a normal-recovery file system, **fsck** performs a multi-pass check.

–p Detect all possible inconsistencies, but correct only those inconsistencies that may be expected to occur from an abnormal system halt. For each corrected inconsistency, one or more lines will be printed identifying the file system and the nature of the correction. Any other inconsistencies will cause the check of that file system to fail. The following 15 inconsistencies (and only those listed) will be corrected for the specified file systems:

1. An inode has an incorrect count of the blocks it uses. The count is corrected.
2. An inode is partially truncated. Partial truncation can occur if the system is abnormally halted while a file is being truncated, leaving the

file claiming more data blocks than its size in bytes would require. The extra blocks are freed.

3. A directory has an incorrect child count. The count is corrected.
4. A directory entry exists for an inode which is unallocated. The directory entry is removed.
5. A directory entry's file name length is incorrect. The length is corrected.
6. An inode is unreferenced (has no directory entries anywhere in the file system). The inode is reconnected in the **lost+found** directory.
7. No **lost+found** directory exists, but an inode needs to be reconnected there. The **lost+found** directory is created.
8. The root directory needs to be expanded in order to make room for a **lost+found** directory entry. The directory is expanded.
9. The **lost+found** directory needs to be expanded in order to make room for a directory entry for an inode being reconnected there. The directory is expanded.
10. An inode's link count is incorrect. The count is corrected.
11. The root control point directory's resource accounting (blocks, inodes) is incorrect. The counts are corrected.
12. A disk allocation region (DAR) has an incorrect free-block bitmap. The bitmap is corrected.
13. A DAR has an incorrect free-inode list. The list is corrected.
14. A DAR has incorrect summary counts of used blocks, inodes or directories. The counts are corrected.
15. The summary counts in the superblock are incorrect. The counts are corrected.

–q Repair the inconsistencies listed under the **–p** option automatically, without asking for user approval. Unlike **–p** however, more serious inconsistencies will not cause **fsck** to fail; the user must still answer the resulting queries.

–y Audit and repair all file system inconsistencies assuming a "yes" response to all questions asked by **fsck**.

CAUTION: Use this option with great care, since it could lead to irreversible changes to the file system.

–n Audit all file system inconsistencies, assuming a "no" response to all questions asked by **fsck**. This option also means that all file systems will be opened with read-only intent.

–x Look at a file system's superblock to see if it is marked mountable. If so, do not check the file system for inconsistencies. If the file system is marked unmountable, check it.

–s Ignore the actual free-block bitmap and unconditionally reconstruct a new one.

–S Conditionally reconstruct the free-block bitmap. A free-block bitmap is reconstructed if and only if the file system is consistent. This option also forces a "no" response to all **fsck** questions.

–D Directories are checked for bad blocks.

–f Fast check: blocks and sizes are checked; the free block bitmap is reconstructed if necessary.

The following options are mutually exclusive, and use of more than one per invocation is not allowed: **–l**, **–y**, **–n**, **–p**, **–q**, and **–S**.

Arguments to fsck Options

If you do not specify which file systems to check on the **fsck** command line, **fsck** will check those having appropriate entries in the **fstab** file. An appropriate entry is one having a nonzero pass number and a "rw" or "ro" mounting status. If the **–p** option was specified, the checking occurs in order of pass number, with those file systems of equal pass number being checked in parallel with each other. Otherwise, checking occurs in order of appearance in **fstab**.

If arguments are specified (the rest of the command line after the option flags), those file systems, and only those file systems, are checked sequentially in the order given.

File systems may be specified as arguments to **fsck** in one of two ways: by the special device file (in **/dev/dsk** or **/dev/rdsk**) containing the file system; or by the directory that **/etc/fstab** indicates will serve as the mount point for the file system.

Checking File Systems

File system checking proceeds without any input from the operator if no errors are discovered. When a fatal inconsistency is discovered, no further checking is done on that file system; the **fsck** program either exits or proceeds to the next specified file system. When an inconsistency is discovered with the **–p** option and that error is one of those listed under **–p**, or when the inconsistency is discovered with the **–y** option, the inconsistency is fixed without operator intervention. Any other discoveries of inconsistencies require the operator to make a decision. The **fsck** program prompts with its recommended action. If you answer **yes**, **fsck** takes the recommended action. In the case of **–p**, no damaging action will be taken without approval. Note, however, that advance approval or disapproval may be given by invoking **fsck** with the **–y** and **–n** options, respectively.

The **fsck** program will refuse to check any file system for which any of the following conditions hold true:

- The file system is mounted (except when you specified the **–n** option, which opens the file system read-only).
- The special file associated with the file system cannot be opened.
- The specified pathname (or its device node associate in **/etc/fstab**) is not a block-special, character-special, or regular file whose size can be determined.

The **fsck** program checks for the following inconsistencies (the term "Bad format" refers to system blocks which do not have the required self-identification information).

- Unreadable or inconsistent superblocks.
- Bad format in superblocks.
- Invalid contents in superblock's reserved area.
- Bad value for superblock's file system size.
- Bad value for superblock's DAR size.
- Bad value for superblock's inode/DAR density.
- Bad value for superblock's default data element size.
- Bad value for superblock's default index element size.
- Bad value for superblock's default directory data element size.
- Bad value for superblock's default directory index element size.
- Bad value for superblock's default first allocation threshold.
- Bad value for superblock's default second allocation threshold.
- Bad format in inode table block.
- Invalid contents in inode's reserved area.
- Files of unknown type.
- Files with bad fragment size.
- Files which are partially truncated.
- Files claiming impossible blocks.
- Files claiming system-area blocks.
- Bad Index-block format.
- Files with incorrect block counts.
- Files claiming already-claimed blocks.
- Unallocated root inode.
- Bad file type for root.
- Incorrect resource limit information in root.
- Incorrect parent directory in root.

- Directories with "holes" (unallocated blocks before end-of-file).
- Bad format in directory blocks.
- Directories with invalid information in reserved areas.
- Directories with empty blocks at end.
- Directories with incorrect child counts.
- Extra directory entries named "." or ".."
- Directory entries with invalid characters in file names: "/" or non-ASCII characters.
- Directory entries which would have too-long pathnames.
- Directory entries which are out of order.
- Directory entries with incorrect entry lengths.
- Directory entries with incorrect file name lengths.
- Extraneous hard links to directories (including cycles in file system name space).
- Extraneous hard links to Symbolic Link files.
- Directory entries to invalid inodes.
- Directory entries to unallocated inodes.
- Files with incorrect space parent.
- Unconnected files or directories.
- Bad or missing **lost+found** directories.
- Bad **lost+found** directory entries.
- Root or **lost+found** directories needing expansion.
- Files with incorrect link counts.
- Incorrect resource allocation counts in control point directories.
- Bad format in DAR blocks.
- Invalid contents in reserved area of DAR blocks.
- Incorrect free-block bitmaps in DARs.
- Incorrect or incomplete free-inode lists in DARs.
- Incorrect DAR summary counts: blocks used, inodes used, directories used.
- Incorrect superblock summary counts.

fsck Output

See Appendix B for a complete list of **fsck** errors.

If the **–p** option is used, **fsck** prints out one or more lines for each inconsistency it corrects, indicating the file system fixed and the error corrected. After successfully checking or correcting a file system, **fsck** prints out the name of the file system, the number of files on it, and the number of free and used blocks.

If you do not specify **–p**, **fsck** is more verbose. It will first print out the name of the file system. If performing a multi-pass check, **fsck** prints a message as it enters each phase of checking a file system. A message is printed for each inconsistency encountered, and the operator is prompted for approval before each correction is attempted. (If the **–y** or **–n** flags are used, **fsck** automatically answers such prompts itself.) When checking is complete for the file system, a message is printed if any corrections were made. Finally, the numbers (used, free and total) of files and blocks are printed.

The **fsck** program attempts to give as much information as possible about any files for which you must make decisions (such as whether to remove it, etc.). At least the following information will always be displayed:

- I-number, which is a number specifying a particular inode on a file system
- Owner's user ID
- Owner's group ID
- File type
- Mode
- Size
- Time of last modification

When possible, the following additional information will be displayed:

- Pathname
- Owner's username
- Owner's group name

End of Chapter

Chapter 10
Managing Terminals and Ports

This chapter tells how to manage your system's ports, including terminal lines and the lines you use for UUCP connections. The DG/UX system provides port services through the Service Access Facility (SAF).

For more information on setting up terminals and keyboards to accommodate your environment, see *Customizing the DG/UX™ System*.

You can manage port services two ways. One way is by using the operations found in **sysadm**'s Port menu, located in the Device menu. The first section of this chapter, "Terminal and Port Operations," covers these operations.

The second way to manage ports is with commands invoked at the shell level. For coverage of the shell commands, see the second part of the chapter, "Expert Information."

Briefly, SAF is controlled by a single process on your system, SAC, the service access controller. The SAC process, which starts when you take the system to run level 2, 3, or 4, is responsible for starting and maintaining any port monitors that you define. A port monitor is in turn responsible for starting and maintaining any port services that you have defined. Figure 10–1 shows the SAF process structure.

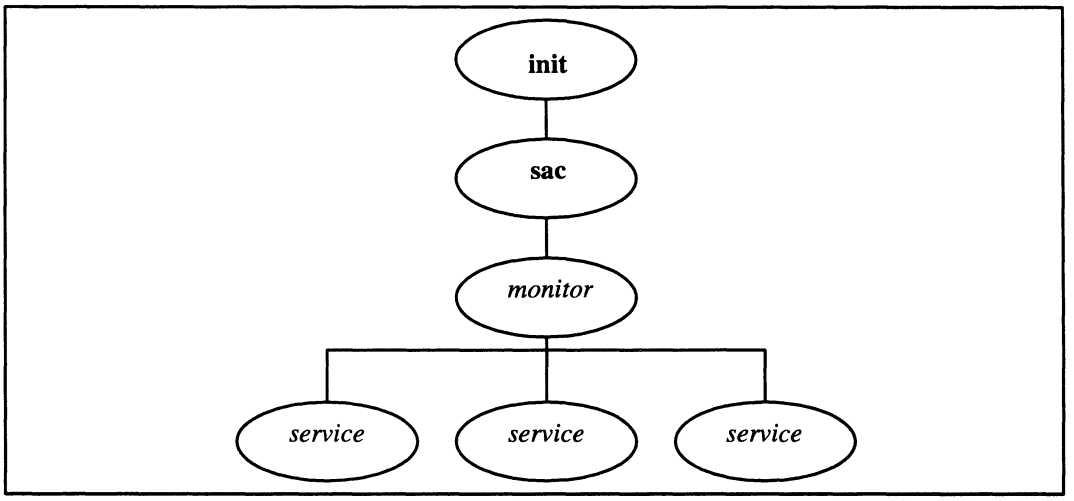

Figure 10–1 SAF Process Structure

For detailed treatment of the Service Access Facility, see the second part of the chapter, "Expert Information."

Terminal and Port Operations

The Port menu provides all the **sysadm** operations necessary for managing terminals and ports. The Terminal menu provides operations for setting up user

terminal lines, sometimes called TTYs. These operations are actually just streamlined versions of the operations in the Port Services menu, offered to make it easier for you to manage user terminal lines.

The Port Monitor menu operations let you set up and maintain port monitors, which are daemons responsible for attaching services to specific ports. The most common function of a port monitor, for example, is to provide login services for user terminals.

The Port Service menu operations let you set up and maintain assignments of port monitors to specific ports. You provide login capabilities to user terminals, for example, by attaching a port monitor to a particular terminal line. The operations in this menu are generalized for setting up all kinds of port services, not just terminal lines.

The following sections discuss the menus and operations in more detail.

Managing Terminals

Although you may manage user terminals (also called TTYs) using the operations in the Port Monitor and Port Service menus, you may find it easier to use the operations in the Terminal menu instead. The operations perform essentially the same functions, but the Terminal menu operations are tailored specifically for managing user terminals.

Adding and Modifying Terminals

The first time you use the `Device -> Port -> Terminal -> Add` operation to add a terminal, the operation checks to see if a **ttymon** port monitor already exists. If one does not, the operation adds and starts a **ttymon** port monitor. The port monitor, whose tag (name) is **ttymon1**, is configured to manage user terminals and provide normal login services. All you have to do is add the terminals with the Add operation.

You should not add a port service on a line that is not connected to a terminal or modem. Lines connected to devices such as printers, the asynchronous port on an uninterruptible power supply unit, or ports used for **mterm**(1) connections may produce noise on the line. Unterminated lines also transmit noise back to the system. Noise on the line can cause the port monitor to consume inordinate amounts of CPU time. Noise on the line can also produce strange errors in port monitor log files. Noise on the line can also produce strange errors in the port monitor log files.

The number of terminals that can be supported by a port monitor is seven less than the value of **HDESLIM**; for example, 1024 – 7 = 1017. In general, however, we recommend that you configure significantly fewer terminals per port monitor.

Additional port monitors can help to share the load incurred by users logging in and logging out. Having multiple port monitors also allows some flexibility in managing user terminals because you can enable and disable a port monitor's terminals as a single unit. Generally, we recommend that you configure port monitors so that you can readily identify devices and terminal lines associated with the port monitor. For

example, you could define a port monitor per cluster box. To add a port monitor, use the operation `Device -> Port -> Port Monitor -> Add`.

The **sysadm** operations Terminal –> Add and Terminal –> Modify present a number of queries. For the Add operation, the defaults are the recommended values for a typical login terminal configuration. For the Modify operation, the defaults are the existing values for the terminal that you are modifying. The prompts for the queries are:

`Controlling port monitor for terminal`
: This prompt appears only if you have more than one **ttymon** port monitor on the system. You need to supply the tag (name) of the port monitor to which this terminal is assigned. The DG/UX system initially has one port monitor, **ttymon1**, for monitoring terminals. This monitor is the default in this query. If you have multiple **ttymon** port monitors on your system, the operation prompts you for the tag of the port monitor to manage the terminal line.

`Tty device(s)`
: Enter the names of the terminal (**tty**) devices from the **/dev** directory, for example, **tty06**, **tty13**.

`TTY definition label`
: Select a terminal definition label. The labels are from the file **/etc/ttydefs**. The number in the label represents line speed. The **M** prefix denotes labels for use with modems, and the **EP** suffix denotes lines with even parity.

`TERM variable`
: Enter the **TERM** variable that will be in effect when the user first logs into the system. The **TERM** variable value that you select indicates the kind of terminal on the line. The **TERM** variable must correspond to an existing **terminfo**(4) database entry. For a complete list of **terminfo** entries, issue this command line at the shell prompt:

```
# ls -CR /usr/share/lib/terminfo/* | more ↵
```

`Disabled response message`
: The disabled-response message is a text string that the system transmits to the terminal when you disable the port. To represent New Line characters and Tab characters in the message, use **\n** and **\t**, respectively.

`Initial state`
: The initial state may be either ENABLED or DISABLED. Users may log in on an enabled terminal. Users cannot log in on a disabled terminal. Users can identify a disabled terminal because the disabled-response message, if defined, appears on it.

Deleting Terminals

When you invoke the Delete operation, you may perform the operation on all terminals or only on terminals that you specify. Deleting a terminal will not terminate a login session currently using the terminal line. When the current user logs off, however, no one will be able to log in over the line.

Listing Terminals

Select the List operation to display attributes of all terminals on your system. The attributes are those set when you added the terminal. See the section above, "Adding and Modifying Terminals," for a discussion of these attributes.

Enabling and Disabling Terminals

Use the Enable operation to allow users to log in using a terminal. The port monitor must also be enabled to allow logging into the system. When you enable a terminal, the system sends the login prompt to the terminal.

Use the Disable operation to disallow users from logging in over the terminal. Disabling a terminal does not terminate any login session currently using the terminal. If you defined a disabled-response message for the terminal, the system sends it to the terminal when you disable it.

Managing Port Monitors

The DG/UX system provides two types of port monitors. The first, **ttymon**, controls access to the system over specified ports, starting programs to provide services such as **login** service to users or applications who attempt to access the system over those ports. The second monitor, **listen**, monitors ports used for TLI (Transport Layer Interface) services. For more information, see *Programming with TCP/IP on the DG/UX™ System*.

The Port Monitor menu provides operations for adding, deleting, modifying, and listing port monitors. It also provides operations for enabling, disabling, starting, and stopping them. By default, SAF starts monitors when it comes up to run level 2, 3, or 4. To change this behavior, edit the second field of the **saf** entry in **/etc/inittab**. See Chapter 3 for more information on the format of **inittab**(4) entries.

Adding and Modifying Port Monitors

To add or modify a port monitor, you need to understand the attributes that define the port monitor. These attributes, discussed below, correspond to the queries in the Add and Modify operations.

`Port Monitor Type`
: The port monitor type may be either **ttymon**, for providing services to terminal users and UUCP callers, or **listen**, for monitoring ports used for TLI communication. If you have written your own port monitor, specify the type here.

`Port Monitor Tag`
: The port monitor tag is a unique name that distinguishes this monitor from others on your system. You may want the name to indicate the ports or terminal lines monitored.

`Command to Start Port Monitor`
: This is the command line that the system should use to start the monitor. Unless you know of options you wish to specify, take the default. See the **ttymon**(1M) manual page or the **listen**(1M) manual for more information on command line invocation.

`Version Number`
: This is the version number of the port monitor. The version number defines the format of the port monitor administrative file.

`Initial Run State`
: The initial run state determines whether or not the system will start the port monitor at boot time. You may choose either STARTED or STOPPED. If you select STOPPED, the port monitor will not start until you start it explicitly with the Start operation or the **sacadm**(1M) or **admportmonitor**(1M) commands.

`Start State`
: The start state determines whether or not the port monitor will accept requests for services when it starts. You may choose either ENABLED or DISABLED. If enabled, the port monitor monitors its ports and accepts connection requests. If disabled, the monitor monitors its ports but does not accept connection requests.

`Restart Count`
: The restart count determines how many times the Service Access Controller (SAC) will attempt to restart the port monitor if it fails. If the port monitor will not start after this number of retries, SAC places it in a FAILED state.

`File Name of Configuration Script`
: A configuration script can initialize the port monitor by pushing STREAMS modules onto the port monitor's stack or by setting environment variables for the port monitor. When the monitor starts, it executes the commands in this script. The configuration script is optional. For more information on configuration scripts, see the **doconfig**(3N) manual page.

`Comment`
: This comment may be any text that you wish to add to describe the port monitor for your own records.

Deleting Port Monitors

The Delete operation removes port monitors. You may specify one or more existing port monitors by tag, or you may perform the operation on all monitors.

Deleting a port monitor terminates all sessions that it started.

Listing Port Monitors

The List operation displays information on all port monitors. An example listing follows:

```
PMTAG     PMTYPE    FLGS RCNT STATUS    COMMAND
tcp       listen    -    3    ENABLED   /usr/lib/saf/listen tcp #listener for tcp
ttymon1   ttymon    -    0    ENABLED   /usr/lib/saf/ttymon     #
```

The # at the end of each line is a comment delimiter.

Enabling and Disabling Port Monitors

You may perform the Enable or Disable operations any time a port monitor is running. By default, the system starts and enables port monitors when the system comes up to run levels 2, 3, and 4.

An enabled port monitor monitors its assigned ports and accepts requests for login connections to the system. A disabled port monitor does not accept requests at any of its ports. Disabling a port monitor terminates all sessions that it started.

Starting and Stopping Port Monitors

If you have not defined your port monitors to start at boot, you need to start them explicitly either with the Start operation or with the **sacadm**(1M) or **admportmonitor**(1M) commands.

Performing the Stop operation for a port monitor does not terminate any connections already established by the port monitor.

Managing Port Services

Adding and Modifying Port Services

After you have added port monitors on your system, you need to assign port services. A port service associates a port (such as a TTY line) with an existing port monitor, describing the nature of the service that the port requires. A port service definition consists of attributes such as the pathname of the device to monitor, the program to run when a user connects on the line, settings determining how communication should occur over the line, and so on. You assign a port service for every port that you want a monitor to manage.

You should not add a port service on a line that is not connected to a terminal or modem. Lines connected to devices such as printers, the asynchronous port on an uninterruptible power supply unit, or ports used for **mterm**(1) connections may produce noise on the line. Unterminated lines also transmit noise back to the system. Noise on the line can cause the port monitor to consume inordinate amounts of CPU time. Noise on the line can also produce strange errors in port monitor log files. Noise on the line can also produce strange errors in the port monitor log files.

Before invoking the Add or Modify operations, familiarize yourself with the queries that they present. Some queries apply to both **ttymon** and to **listen** port services while others apply only to one or the other.

Queries for Adding or Modifying ttymon and listen Services

`Controlling port monitor for service`

This is the tag (name) of the port monitor that will monitor the port. You should have already added this port monitor with the **sysadm** operation `Device -> Port -> Port Monitor -> Add`.

`Port service tag`

This is a name representing this service definition. This tag may be any name that you choose. The name should be unique for the given port monitor. For port services added with the Add operation of the Terminal menu, the tag is the name of the TTY, for example, `tty02`.

`Service Userid`

Every service runs as a process on your system. The user ID you select determines which user will own the process. The user ID must already exist. Typically, the user ID is **root**.

`Create utmp entry`

You may choose whether or not the system creates an **/etc/utmp** entry for the connecting user. Several system programs, including **who**, **write**, and **login**, depend on the **utmp** file for user information.

`File name of configuration script`

You may specify a configuration script to modify the behavior of the port monitor. This configuration script affects only the service on this port. The port monitor copies the script to its own configuration script when it starts. For more information on configuration scripts, see the **doconfig**(3N) manual page.

`Comment`

This comment may be any text that you choose. You may use the comment field to contain a description of the port service for your records.

`Initial state`

The initial state may be either ENABLED or DISABLED. An enabled line accepts requests for connections to the service. A typical login line should be enabled to allow users to log into the system.

A disabled line refuses requests. When you disable a line, the system transmits a disabled-response message on the line. You enter the disabled-response message in the "Disabled response message" query, below.

`Version number`

This field is the version number of the port monitor.

Queries for Adding or Modifying ttymon Services

`Path name of terminal device`

Specify the pathname of the port or terminal. This path should indicate the port's entry in the **/dev** directory, for example, **`/dev/tty04`**.

`TTY definition label`

This label corresponds to an entry in the **/etc/ttydefs** file. The TTY definition determines terminal I/O characteristics for the line, such as line speed and whether or not the line is for a modem. In the TTY definitions provided with the system, the number in the label indicates line speed. The prefix **`M`** indicates that the definition is intended for modems, and the suffix **`EP`** indicates that the line is even parity.

`Service command`

The service command is the shell command line that the monitor should invoke for a calling user when the line is enabled. For a typical login line, the service is **/usr/bin/login**.

`Hangup`

This attribute causes the system to "hang up" the line by setting line speed to zero before initializing the line. This feature is useful if the line is connected to a modem.

`Connect on carrier`

This feature causes the port monitor to invoke the service as soon as it receives a carrier indication. Use this feature only if you are sure of the baud rate of the incoming caller and you know that the incoming caller does not require a prompt to begin the session.

`Bidirectional`

A bidirectional line allows local users and programs to call out on the line and outside users to call in on the line. For a normal terminal line, you do not want to set the line for bidirectional use. If the line is connected to a modem used for dialing out as well as for receiving calls, you want to set the line for bidirectional use.

`Wait-read value`

This value determines how many New Line characters the port monitor must receive before it sends the prompt out on the line. If you specify 0, the port monitor waits for any character. If you specify "none," **ttymon** sends a prompt without waiting for characters.

`Timeout`

This value determines how long the line may be inactive at the login prompt before the port monitor terminates the session. An inactive line is one over which the port monitor detects no transmission of characters. For no time-out period, specify zero seconds.

`Prompt message`

The port monitor transmits the prompt message when it places a port in the ENABLED state. To represent New Line characters and Tab characters in the message, use `\n` and `\t`, respectively.

`Modules to be pushed`

Enter the additional STREAMS modules that you want pushed to handle the line. By default, lines connected to a **syac** already have STREAMS modules **ldterm** and **ttcompat** pushed. For more information on STREAMS, see *Programming in the DG/UX™ Kernel Environment* and *UNIX System V Release 4: Programmer's Guide: STREAMS*.

`Disabled response message`

Enter the message that you want the port monitor to transmit when callers attempt to use a disabled line. To represent New Line characters and Tab characters in the message, use **\n** and **\t**, respectively. The default is no message.

Queries for Adding or Modifying listen Services

`Service type`

You may select either of two values for this field:

`Spawn a service`

Select this value to cause the port monitor to invoke the service that you specify in the next field, `Service command or STREAMS pipe`, whenever a connection request is received.

`Pass file descriptor to standing server`
: Select this value to cause the port monitor to pass the file descriptor for this connection through the pipe you specify in the next query, `Service command or STREAMS pipe`, whenever a connection request is received. The pipe is connected to a standing server, a server that is currently running.

`Service command or STREAMS pipe`
: If you selected `Spawn a service` for the service type, enter the service to be started. If you selected `Pass file descriptor to standing server` in the previous query, enter the name of the STREAMS pipe to be used for passing file descriptors.

`Modules to be pushed`
: Enter the names of the STREAMS modules to be pushed. After popping all modules already on the stream, the operation pushes the specified modules in the order in which you enter them.

`Server's private address`
: Enter the address that **listen** should monitor. **listen** will dispatch calls at this address to the designated service. The address must be unique.

Deleting Port Services

Use the Delete operation to remove port service definitions. You may perform the operation on all port services or only on those specified.

Listing Port Services

Use the List operation to display port service definitions. The operation displays definitions for all port services. The display is generated by **pmadm –l**, discussed in "Expert Information."

Enabling and Disabling Port Services

Use the Enable operation to make a port service available. The service is available provided the port monitor is also enabled. For typical terminal lines, this means that users can log into the system.

Use the Disable operation to make a port service unavailable. When the port monitor detects an attempt to access a disabled service, it transmits the disabled-response message defined for the port service.

Expert Information

This section provides technical detail about SAF (Service Access Facility). The preceding portion of the chapter tells how to use the **sysadm** operations to manage port services; this section describes the underlying functionality.

The Service Access Facility generalizes the procedures for service access so that login access on the local system and network access to local services are managed in essentially similar ways.

Systems may access services using a variety of port monitors, including port monitors written expressly for a user's application. The section on **listen** describes the administrative tasks required to provide access to services over any network that conforms to the Transport Layer Interface (TLI) protocol. For more information on TLI, see *Programming with TCP/IP on the DG/UX™ System*.

The "Expert Information" section is divided into six subsections. The first, "Overview of the Service Access Facility," describes

- the Service Access Facility (SAF),
- its directory structure,
- its controlling program (the Service Access Controller or SAC),
- the configuration files that may be used to change the environment under which SAC, port monitors, and services operate, and
- the SAF's two administrative files.

The second section, "Port Monitor Management," describes the SAC administrative command **sacadm** and how it is used to manage the information in the SAC administrative file. This includes:

- printing port monitor status information
- adding a port monitor to the system
- enabling or disabling a port monitor
- starting or stopping a port monitor, and
- removing a port monitor from the system.

This section also describes the per-system and per-port monitor configuration scripts used to modify the SAC and port monitor environments and shows how to print, install, and modify each type of script.

A summary of the commands used to administer a port monitor is included at the end of the section.

The third section, "Service Management," follows the same outline. It describes the port monitor administrative command, **pmadm**. **pmadm** manages the information in the port monitors' administrative files. These files, one for each port monitor, include administrative and state information for each port and information about the service each port invokes. The **pmadm** command allows the system administrator to

- print information derived from the administrative file
- add and remove services
- enable and disable services
- print, install, and change per-service configuration scripts

A command summary is included at the end.

The fourth section, "The Port Monitor ttymon" describes what **ttymon** does and shows how to perform some of the administrative tasks described in the sections on port monitor management and service management—this time from the point of view of **ttymon**. These tasks include:

- listing the **ttymon** port monitors that have been configured
- listing the services that have been configured under a given **ttymon** port monitor
- enabling and disabling **ttymon** ports and services
- adding and removing a **ttymon** port monitor
- adding a service under a **ttymon** port monitor

The fifth section, "Terminal Line Settings," covers the **ttydefs** file, the database file for system TTY terminal information. The **ttymon** port monitor uses the file and the system administrator modifies it in the process of **ttymon** administration. Setting terminal modes and line speeds includes:

- printing terminal line setting information
- modifying terminal line settings
- setting up hunt sequences
- adding and removing terminal line settings for a terminal
- setting terminal options available using the **stty** command

A summary of commands used in **ttymon** and terminal line settings administration is included at the end of the section.

The sixth section, "The Port Monitor listen," describes what **listen** does and how you perform **listen**-related administrative tasks. These include:

- listing configured **listen** port monitors
- listing services available through a given **listen** port monitor
- enabling and disabling **listen** ports and services
- adding and removing **listen** port monitors
- adding a **listen** service

Tasks associated with port and service administration may be performed using either the operations in the **sysadm** utility or shell commands entered on the command line.

The shell commands for managing port monitors and port services are described in "Port Monitor Management" and "Service Management."

Overview of the Service Access Facility

The Service Access Facility (SAF) provides a mechanism for uniform access to services. From the system administrator's point of view, the main components of this generalized access procedure are the commands for installing, configuring, and maintaining port monitors and services and the administrative or database files in which port monitor and service information is stored.

From the point of view of the Service Access Facility, a service is a process that is started. There are no restrictions on the functions a service may provide.

The Service Access Facility consists of a controlling process, the Service Access Controller (SAC), and two administrative levels corresponding to two levels in the supporting directory structure. The top administrative level is concerned with port monitor administration, the lower level with service administration.

From an administrative point of view, the Service Access Facility consists of the following components, each of which is described in this section:

- The Service Access Controller
- A per-system configuration script
- The SAC administrative file
- The SAC administrative command **sacadm** or **admportmonitor**(1M)
- Port monitors
- Optional per-port monitor configuration scripts
- An administrative file for each port monitor
- The administrative command **pmadm** or **admportservice**(1M)
- Optional per-service configuration scripts

The Service Access Controller, the administrative files, and the per-system, per-port monitor, and per-service configuration files are described in this section. The administrative command **sacadm** is described under "Port Monitor Management" and the administrative command **pmadm** is described under "Service Management." The **sysadm** utility performs SAC administrative tasks with the **admportmonitor**(1M) command and port monitor-specific administrative tasks with the **admportservice**(1M) command.

The Service Access Controller

The Service Access Controller is the overseer of the server system. It is the Service Access Facility's controlling process. SAC is started by **init**(1M) by means of an entry in **/etc/inittab**. Its function is to maintain the port monitors on the system in the state specified by the system administrator. These states include: STARTING, ENABLED, DISABLED, STOPPING, NOT RUNNING, and FAILED. (A port monitor enters the FAILED state if SAC cannot start it after a specified number of tries.)

The administrative command **sacadm** is used to tell SAC to change the state of a port monitor. **sacadm** can also be used to add or remove a port monitor from SAC supervision and to list information about port monitors known to SAC. The **admportmonitor**(1M) command also provides these services.

SAC's administrative file contains a unique tag for each port monitor known to SAC and the pathname of the command used to start each port monitor.

SAC performs three main functions. Briefly:

- It customizes its own environment.
- It starts the appropriate port monitors.
- It polls its port monitors and initiates recovery procedures when necessary.

During initialization, SAC customizes its own environment by invoking the per-system configuration script. Next, it reads its administrative file to determine which port monitors are to be started. For each port monitor it starts, it interprets the port monitor's configuration script, if one exists. Finally the port monitors specified in the administrative file (for example, **ttymon**) are started.

Once the port monitors are running, SAC polls them periodically for status information. The **sac**(1M) command line option, **–t**, allows the system administrator to control polling frequency. By examining the **saf** entry in the **/etc/inittab** file, you can see that the default is to poll every 45 seconds. When the port monitor gets a request for status from SAC, it must respond with a message containing its current state (for example, ENABLED). If SAC does not receive a response, it assumes the port monitor is not running. If a port monitor that should be running has stopped, SAC assumes it has failed and takes appropriate recovery action.

SAC will restart a failed port monitor if a nonzero restart count was specified for the port monitor when it was created (see the **sacadm** command, described under "Port Monitor Management," below, and in the **sacadm**(1M) manual page.

SAC is the administrative point of control for all port monitors (and therefore for all ports on the system). The administrative commands **sacadm**(1M) and **pmadm**(1M) pass requests to SAC, which in turn communicates with the port monitors. These requests include enabling a disabled port monitor so that it begins accepting service requests on its ports; starting port monitors that were previously killed; and listing the current state of all port monitors on the system.

The Per-System Configuration File

The prototype per-system configuration file, **/etc/saf/_sysconfig.proto**, is delivered empty. The system administrator can customize the environment for all services on the system by writing a command script in the interpreted language described in the *Managing TCP/IP on the DG/UX™ System* and in the **doconfig**(3N) manual page and placing this command script in the **_sysconfig** file. The per-system configuration script is interpreted by the Service Access Controller when SAC is started. SAC is started when the system enters multiuser mode.

Per-Port Monitor Configuration Scripts

Per-port monitor configuration scripts (**/etc/saf/***pmtag*/**_config**) are optional. They allow the system administrator to customize the environment for any given port monitor and for the services that are available through the specific collection of access points for which that port monitor is responsible. Per-port monitor configuration scripts are written in the same language used for per-system configuration scripts.

The per-port monitor configuration script is interpreted when the port monitor is started. The port monitor is started by the Service Access Controller after SAC has itself been started and after it has run its own configuration script, **/etc/saf/_sysconfig**.

The per-port monitor configuration script may override defaults provided by the per-system configuration script.

Per-Service Configuration Scripts

Per-service configuration files allow the system administrator to customize the environment for a specific service. For example, a service may require special privileges that are not available to the general user. Using the language described in the **doconfig**(3N) manual page, the system administrator can write a script that will grant or limit such special privileges to a particular service offered through a particular port monitor.

The per-service configuration may override defaults provided by higher-level configuration scripts. For example, the per-service configuration script may specify a set of STREAMS modules other than the default set.

The SAC Administrative File

SAC's administrative file contains information about all the port monitors for which SAC is responsible. The prototype SAC administrative file, **/etc/saf/_sactab.proto**, contains a comment line containing the version number of the Service Access Controller and an entry for the **tcp** port monitor. The system administrator adds port monitors to the system by using the administrative command **sacadm** with the **–a** option to make entries in SAC's administrative file. **sacadm** is also used to remove entries from SAC's administrative file. As an alternative, you can use the **admportmonitor**(1M) command or **sysadm**'s `Device -> Port -> Port Monitor` operations to perform the same functions.

NOTE: We recommend that you do not change the SAC administrative file except through the **sacadm**(1M), **admportmonitor**(1M), or **sysadm**(1M) utilities. If you choose to modify the SAC administrative file, do not change it while SAF is running.

Each entry in SAC's administrative file contains the following information:

PMTAG

A unique tag that identifies a particular port monitor. The system administrator is responsible for naming a port monitor. This tag is then used by the Service Access Controller (SAC) to identify the port monitor for all administrative purposes.

PMTAG may consist of up to 14 alphanumeric characters.

PMTYPE

The type of the port monitor. In addition to its unique tag, each port monitor has a type designator. The type designator identifies a group of port monitors that are different invocations of the same entity. **ttymon** and **listen** are examples of valid port monitor types. The type designator is used to facilitate the administration of groups of related port monitors. Without a type designator, the system administrator has no way of knowing which port monitor tags correspond to port monitors of the same type.

PMTYPE may consist of up to 14 alphanumeric characters.

FLGS The flags that are currently defined are:

d When started, do not enable the port monitor

x Do not start the port monitor.

If no flag is specified, the default action is taken. By default a port monitor is started and enabled.

RCNT The number of times a port monitor may fail before being placed in a failed state. Once a port monitor enters the failed state, SAC will not try to restart it. If a count is not specified when the entry is created, this field is set to **0**. A restart count of **0** indicates that the port monitor is not to be restarted when it fails.

COMMAND

The command line that starts the port monitor. The first component of the string, the command itself, must be a full pathname.

The example below shows the contents of a sample SAC administrative file as listed by the **sacadm** command.

```
# sacadm -l ↵
PMTAG     PMTYPE    FLGS RCNT STATUS    COMMAND
tcp       listen    -    3    ENABLED   /usr/lib/saf/listen tcp #listener for tcp
ttymon1   ttymon    -    0    ENABLED   /usr/lib/saf/ttymon #
```

The # character at the end of each line is a comment delimiter.

The Port Monitor Administrative File

Each port monitor has its own administrative file, **/etc/saf/***pmtag***/_pmtab**. The **pmadm**(1M) command is used to add, remove, or modify entries in this file. The **admportservice**(1M) command also performs these functions. Each time a change is made, the corresponding port monitor is told to reread its administrative file.

NOTE: We recommend that you do not change a port monitor administrative file except through the **pmadm**(1M) or **admportservice**(1M) commands. If you choose to modify a port monitor administrative file, do not change it while the port monitor is running.

Each entry in a port monitor's administrative file defines how the port monitor should treat a specific port and what service is to be invoked on that port. Some fields must be present for all types of port monitors. Each entry must include a service tag to identify the service uniquely and an identity to be assigned to the service when it is started (for example, **root**). The combination of a service tag and a port monitor tag uniquely define an instance of a service. The same service tag may be used to identify a service under a different port monitor. The record must also contain port monitor specific data such as the prompt string which is meaningful to **ttymon**. In general, each type of port monitor provides a command that takes the necessary port monitor-specific data as arguments and outputs these data in a form suitable for storage in the file. The **ttyadm**(1M) command does this for **ttymon** and **nlsadmin**(1M) does it for **listen**.

Each entry in the port monitor administrative file must contain the following information.

SVCTAG
A unique tag that identifies a service. This tag is unique only for the port monitor through which the service is available. Other port monitors may offer the same or other services with the same tag. A service requires both a port monitor tag and a service tag to identify it uniquely.

SVCTAG may consist of up to 14 alphanumeric characters.

FLGS Flags with the following meanings may currently be included in this field:

x Do not enable this port. By default the port is enabled.

u Create a **utmp** entry for this service. By default no **utmp** entry is created for the service.

Note that port monitors may ignore the **u** flag if creating a **utmp** entry for the service is not appropriate to the manner in which the service is to be invoked. Some services may not start properly unless **utmp** entries have been created for them (for example, **login**).

ID The identity under which the service is to be started. The identity has the form of an existing login name.

PMSPECIFIC
Examples of port monitor-specific information are addresses, the name of a process to execute, or the name of a STREAMS pipe through which to pass a connection. See the **ttyadm**(1M) manual page for information about additional **PMSPECIFIC** items.

COMMENT
A comment associated with the service entry.

Each port monitor administrative file must contain one special comment of the form:

```
# VERSION=value
```

where *value* is an integer that represents the port monitor's version number. The version number defines the format of the port monitor administrative file. This comment line is created automatically when a port monitor is added to the system. It appears on a line by itself, before the service entries.

The example below shows the contents of a sample **ttymon** administrative file as listed by the **pmadm** command.

```
# pmadm -l -p ttymon1 ↵
PMTAG   PMTYPE SVCTAG  FLGS ID   <PMSPECIFIC>
ttymon1 ttymon ttymon1 u    root /dev/tty00 bhr 8 /usr/bin/login 60 \
       M1200 - login:  -  #
```

The # character at the end of each line is a comment delimiter.

Note that everything in the PMSPECIFIC column is specific to a **ttymon** port monitor. The listing for a **listen** administrative file, for example, will contain a different set of entries in this column. Port-monitor specific information is formatted by the port monitor's administrative command, in this case **ttyadm**. The **ttyadm** command is included as part of the **pmadm** command when it is used with the **–a** option. See "Adding a Service," under "Service Management," below.

To maintain the integrity of the system, it is strongly recommended that changes in the SAC and port monitor administrative files be made with the **sacadm** and **pmadm** commands, not by editing the files. SAC does not recognize changes in some of the fields in these files unless they are made using the appropriate administrative command. Editing the file directly can lead to unexpected results.

Port Monitor Management

The Service Access Facility administrative model is hierarchical. The highest level is concerned with port monitor administration. The lower level is concerned with service administration and is discussed under "Service Management," below. At the level of port monitor administration, port monitors may be added, removed, started, stopped, enabled, or disabled. Other functions performed at this level include requesting port monitor status information, replacing a per-system configuration file, installing or replacing a per-port monitor configuration file, and requesting that a port monitor read its administrative file.

Configuration files are described under "Printing, Installing, and Replacing Configuration Scripts." Requesting that a port monitor read its administrative file is under "Reading the Administrative Files."

The SAC Administrative Command sacadm

sacadm is the administrative command for the upper level of the Service Access Facility hierarchy, that is, for port monitor administration (see the manual page **sacadm**(1M). Under the Service Access Facility, port monitors are administered by using the **sacadm** command to make changes in SAC's administrative file. **sacadm** performs the functions listed below. Each function is discussed in one of the following sections.

- Print requested port monitor information from the SAC administrative file.
- Add or remove a port monitor.
- Enable or disable a port monitor.
- Start or stop a port monitor.
- Install or replace a per-system configuration script.
- Install or replace a per-port monitor configuration script.
- Ask SAC to reread its administrative file.

Printing Port Monitor Status Information

Unless the system administrator already knows the type of a port monitor, it may be necessary to use the most general form of the command (**sacadm –l**) to find out what the valid type and tag names are.

```
sacadm -L [ -p pmtag | -t type ]
sacadm -l [ -p pmtag | -t type ]
```

pmtag is the tag associated with the port monitor that is being listed. *type* specifies the port monitor type, for example, **listen**.

The command options function as follows:

-l	By itself, the **–l** option lists status information for all services on the system.
-l -p *pmtag*	Lists status information for all services available through port monitor *pmtag*.
-l -t *type*	Lists status information for all services available through port monitors of type *type*.

Other combinations of options with **–l** are invalid.

The **–L** option is identical to the **–l** option except that its output is printed in a condensed format.

Options that request information write the requested information to the standard output. A request for information using the **–l** option prints column headers and aligns the information under the appropriate headings. A request for information in the condensed format using the **–L** option prints the information in colon-separated fields. If the **–l** option is used, empty fields are indicated by a hyphen. If the **–L** option is used, empty fields are indicated by two successive colons.

The following sample output shows the differences between some of the options described above.

```
# sacadm -l ↵
PMTAG   PMTYPE  FLGS RCNT STATUS  COMMAND
tcp     listen  -    3    ENABLED /usr/lib/saf/listen tcp #listener for tcp
ttymon1 ttymon  -    0    ENABLED /usr/lib/saf/ttymon #
```

This is the most general form of the list option.

The STATUS field indicates the monitor's current state, whether disabled or enabled. Entries in the FLGS field do not change when a **ttymon** monitor changes from enabled to disabled or vice-versa. The FLGS field conveys information about the state in which a port monitor starts, not about its current state. For example, the **d** flag indicates that the port monitor goes immediately to DISABLED state when it is started.

The following command lists status information only for port monitor **tcp**:

```
# sacadm -l -p tcp ↵
PMTAG   PMTYPE   FLGS RCNT STATUS    COMMAND
tcp     listen   -    3    ENABLED   /usr/lib/saf/listen tcp #listener for tcp
```

The same command using **-L** instead of **-l** will produce:

```
# sacadm -L -p tcp ↵
tcp:listen::3:DISABLED:/usr/lib/saf/listen -m tcp tcp # tcp listener
```

The following command lists status information for all port monitors whose type is **ttymon**:

```
# sacadm -l -t ttymon ↵
PMTAG     PMTYPE     FLGS RCNT STATUS     COMMAND
ttymon1   ttymon     -    0    ENABLED    /usr/lib/saf/ttymon #
```

Adding a Port Monitor

sacadm -a -p *pmtag* **-t** *type* **-c** *"cmd"* **-v** *ver* [**-f dx**] [**-n** *count*] \
[**-y** *"comment"*] [**-z** *script*]

The **sacadm** command with the **-a** option creates new instances of a port monitor. Because of the complexity of the options and arguments that follow the **-a** option, it may be advisable for the system administrator to use a command script or the **sysadm** operation Device –> Port –> Port Monitor –> Add to add port monitors.

When **sacadm** creates a port monitor, it creates the supporting directory structure in **/etc/saf** and **/var/saf** for the new port monitor *pmtag* and the port monitor administrative file. It also adds an entry for the new port monitor to the SAC's administrative file.

The options following the **-a** option have the following meanings:

- The **-t** option is followed by the port monitor type. The type can be either **ttymon** or **listen**.
- The **-c** option is followed by a command enclosed in double quotes. This is the command SAC executes to start the port monitor.
- The **-v** option is followed by the version number of the port monitor. The version number may be given to **sacadm** by the port monitor's special administrative command, as an argument to the **-v** option. For example:

  ```
  -v `ttyadm -V`
  ```

The port monitor-specific command is **ttyadm** for **ttymon** and **nlsadmin** for **listen** (see the manual pages **ttyadm**(1M) and **nlsadmin**(1M)). The version stamp of the port monitor is known by the command and is returned when the port monitor administrative command is invoked with the **–V** option. The version number is added to the new administrative file as a comment line of the form

```
# VERSION=value
```

where *value* is an integer that represents the port monitor version number. The version number defines the file format. It provides a means of synchronizing software releases of port monitors with their properly formatted administrative files.

- The **–f** option specifies one or both of the two flags **d** and **x**. The flags have the following meanings:

  d Do not enable the port monitor

  x Do not start the port monitor

  If the **–f** option is not included in the command line no flags are set and the default conditions prevail. By default a port monitor is started and enabled.

- The **–n** option sets the restart count to *count*. The restart determines how many times SAC will attempt to start the port monitor before placing it in the FAILED state. If a restart count is not specified when adding a port monitor, *count* is set to **0**. A count of **0** indicates that the port monitor is not to be restarted if it fails.

- The **–y** option includes *"comment"* in the SAC administrative file entry for the port monitor being added.

- The **–z** option names a file whose contents are installed as the per-port monitor configuration script, **_config**.

The command line below adds a port monitor of type **listen** (for TCP/IP).

```
# sacadm -a -p tcp -t listen -c "/usr/lib/saf/listen -m tcp tcp" \
   -v `nlsadmin -V` )
```

The command line in the following figure adds a port monitor of type **ttymon**.

```
# sacadm -a -p ttymon1 -t ttymon -c "/usr/lib/saf/ttymon" -v `ttymon -V` )
```

Enabling, Disabling, Starting, and Stopping a Port Monitor

The commands to enable, disable, start, and stop a port monitor use the following syntax:

```
sacadm -e -p pmtag
sacadm -d -p pmtag
sacadm -s -p pmtag
sacadm -k -p pmtag
```

The **–e** option enables a port monitor. SAC sends an enable message to the port monitor.

The **–d** option disables a port monitor. SAC sends a disable message to the port monitor.

The **–s** option starts a port monitor.

The **–k** option kills a port monitor. SAC sends the signal SIGTERM to the port monitor.

Removing a Port Monitor

sacadm -r -p *pmtag*

Invoke **sacadm** with the **–r** option to remove port monitor *pmtag* from the system. The port monitor entry is removed from SAC's administrative file and SAC rereads the file. If the removed port monitor is not running, no further action is taken. If the removed port monitor is running, the Service Access Controller sends it a SIGTERM signal to indicate that it should shut down. Note that the port monitor's directory structure remains intact but is no longer referenced by anything.

Printing, Installing, and Replacing Configuration Scripts

Per-system and per-port monitor configuration scripts are administered using **sacadm**; per-service configuration scripts are administered using **pmadm** and are described under "Service Management" below. Per-system and per-port monitor configuration scripts allow the system administrator to modify the system and port monitor environments. They are written in the interpreted language described in the manual page for **doconfig**(3N). Sample configuration scripts are shown below.

The per-system configuration script, **_sysconfig**, is interpreted when SAC is starting. A port monitor's per-port monitor configuration script is interpreted by SAC just before SAC starts the port monitor.

Per-system and per-port monitor configuration scripts may be printed by any user on the system. Only the system administrator may install or replace them.

Per-System Configuration Scripts

sacadm -G [**-z** *script*]

The per-system configuration script **/etc/saf/_sysconfig** customizes the environment for all services on the system. When it starts up, the Service Access Controller interprets the per-system configuration script, using the **doconfig** library routine. The prototype **_sysconfig** file, **/etc/saf/_sysconfig.proto**, contains only a comment line.

The **–G** option is used to print or replace the per-system configuration script. The **–G** option by itself prints the per-system configuration script to the screen. The **–G** option in combination with the **–z** option replaces **/etc/saf/_sysconfig** with the contents of the file *script*. Other combinations of options with the **–G** option are invalid.

The sample **_sysconfig** file below sets the time zone variable, TZ.

```
assign TZ=EST5EDT          # set TZ
runwait echo SAC is starting > /dev/console
```

The **–z** option is also used with the **–a** option to specify the contents of the per-port monitor configuration file when a port monitor is created.

Per-Port Monitor Configuration Scripts

sacadm **-g** **-p** *pmtag* [**-z** *script*]

The per-port monitor configuration script **/etc/saf/***pmtag***/_config** customizes the environment for services that are available through the specific collection of access points for which port monitor *pmtag* is responsible. When SAC starts a port monitor, the per-port monitor configuration script is interpreted, if it exists, using the **doconfig**(3N) library routine.

The **–g** option is used to print, install, or replace a per-port monitor configuration script. A **–g** option requires a **–p** option. The **–g** option with only a **–p** option prints the per-port monitor configuration script for port monitor *pmtag*. The **–g** option with the **–p** option and a **–z** option installs the contents of file *script* as the per-port monitor configuration script for port monitor *pmtag*, or, if **/etc/saf/***pmtag***/_config** exists, it replaces **_config** with the contents of *script*. Other combinations of options with **–g** are invalid.

In the following sample **_config** file, the command **/usr/bin/daemon** is assumed to start a daemon process that builds and holds together a STREAMS multiplexor. By installing this configuration script, the command can be executed just before starting the port monitor that requires it.

```
run /usr/bin/daemon
# build a STREAMS multiplexor.
runwait echo $PMTAG is starting > /dev/console
```

Reading the Administrative Files

sacadm -x [**-p** *pmtag*]

When changes are made to SAC's administrative file, SAC needs to be notified of the change. When changes are made to a port monitor's administrative files, the port monitor needs to be notified. When **sacadm** and **pmadm** are used to make changes, this notification takes place automatically. If the files are edited by the system administrator, SAC and the port monitors are not notified. In this case, **sacadm** must be called with the **–x** option to notify SAC or port monitor of the changes.

The **–x** option tells SAC to update its internal copy of the information in the SAC administrative file. **sacadm** with the **–x** and **–p** options causes SAC to send a READ message to the designated port monitor.

System administrators are advised against editing these files directly.

Appendix C contains a reference table that you may find useful for managing port monitors.

Service Management

The top level of the Service Access Facility is concerned with port monitor administration and is discussed in "Port Monitor Management" above. The lower level is concerned with service administration and is discussed in this section.

At this level there are two distinct administrative functions. The first is the administration of the port itself. The information needed to administer a terminal port will be found on the manual page for **ttymon**'s port monitor-specific command, **ttyadm**(1M). The information needed to administer a network address monitored by a **listen** port monitor will be found on the manual page for **listen**'s port monitor-specific command, **nlsadmin**(1M).

The second is the administration of the service associated with a port. By definition, there is one and only one service associated with a port. All ports on the system are peers and their services are administered through the same command interface, the Service Access Facility's administrative command **pmadm**(1M). At the level of service administration, services may be added, removed, enabled, and disabled. Other functions performed at this level include installing or replacing a per-service configuration script and requesting service status information.

The Port Monitor Administrative Command pmadm

pmadm is the administrative command for the lower level of the Service Access Facility hierarchy, that is, for service administration. A port may have only one service associated with it although the same service may be available through more than one port. **pmadm** performs the following functions:

- print service status information from the port monitor's administrative file
- add or remove a service
- enable or disable a service
- install or replace a per-service configuration script

Note that in order to identify an instance of a service uniquely, the **pmadm** command must identify both the service (**–s**) and the port monitor or port monitors through which the service is available (**–p** or **–t**).

Printing Service Status Information

```
pmadm -l  [ -t type | -p pmtag ] [ -s svctag ]
pmadm -L  [ -t type | -p pmtag ] [ -s svctag ]
```

The **–l** and **–L** options request service status information. They may be invoked by any user on the system. Used either alone or with the options described below they provide a filter for extracting information in several different groupings.

–l	By itself, the **–l** option lists status information for all services on the system.
–l –p *pmtag*	Lists status information for all services available through port monitor *pmtag*.
–l –s *svctag*	Lists status information for all services with the tag *svctag* available through any port monitor on the system.
–l –p *pmtag* **–s** *svctag*	Lists status information for service *svctag* available through port monitor *pmtag*.

–l –t *type*	Lists status information for all services available through port monitors of type *type*.
–l –t *type* **–s** *svctag*	Lists status information for all services with the tag *svctag* offered through a port monitor of type *type*.

Other combinations of options with **–l** are invalid.

The **–L** option is identical to the **–l** option except that output is printed in a condensed format.

Options that request information write the requested information to the standard output. A request for information using the **–l** option prints column headers and aligns the information under the appropriate headings. A request for information in the condensed format using the **–L** option prints the information in colon-separated fields. If the **–l** option is used, empty fields are indicated by a hyphen. If the **–L** option is used, empty fields are indicated by two successive colons.

The example below shows a sample of **–l** output. The lines have been broken for readability.

```
PMTAG PMTYPE SVCTAG FLGS ID    PMSPECIFIC
tmon3 ttymon 6      ux   root   /dev/tty06 - - /usr/bin/login - 9600 - login: - \
                                #tty6
tmon3 ttymon 7      ux   root   /dev/tty07 - - /usr/bin/login - 9600 - login: - \
                                #tty7
tmon3 ttymon 8      ux   root   /dev/tty08 - - /usr/bin/login - 9600 - login: - \
                                #tty8
tmon3 ttymon 9      ux   root   /dev/tty09 - - /usr/bin/login - 9600 - login: - \
                                #tty9
tmon1 ttymon 1      ux   root   /dev/tty01 - - /usr/bin/login - 9600 - login: - \
                                #tty1
tmon1 ttymon 2      ux   root   /dev/tty02 - - /usr/bin/login - 9600 - login: - \
                                #tty2
tmon1 ttymon 3      ux   root   /dev/tty03 - - /usr/bin/login - 9600 - login: - \
                                #tty3
tmon1 ttymon 4      ux   root   /dev/tty04 - - /usr/bin/login - 9600 - login: - \
                                #tty4
tcp   listen 101    -    listen - c - /usr/lib/uucp/uucico -r0 -unuucp -iTLI \
                                #UUCP access direct to server
tcp   listen 102    -    listen - c - /usr/tcp/lib/ttysrv -d -n ntty,tirdwr,1d0 \
                                #UUCP access to server via login
tcp   listen 1000   -    listen - c - /usr/tcp/lib/tstserver #TP TEST SERVER
tcp   listen 0      -    root   - c - sfbul.serve c - /usr/lib/saf/nlps_server
```

Adding a Service

pmadm -a [**-p** *pmtag* | **-t** *type*] **-s** *svctag* **-i** *id* **-m** *"pmspecific"* **-v** *ver* ****
    [**-f xu**] [**-y** *"comment"*] [**-z** *script*]

The **–a** option adds a service by making an entry for the new service in the port monitor's administrative file. It is important to be aware that a service implies a port and that there is a one-to-one mapping between ports and instances of services. Because of the complexity of the options and arguments that follow the **–a** option, it may be advisable to use a command script or **sysadm** to add services. If you use **sysadm**, select the operation Device –> Port –> Port Service –> Add.

The following paragraphs describe the components of the command line for adding a service.

-p *pmtag* Specify the tag for the port monitor. This option causes **pmadm** to add the service to the port monitor designated as *pmtag*. This is a name that the administrator selects.

-s *svctag* Specify the tag for the port service. This is a name that the administrator selects.

-t *pmtype* Specify a group of port monitors by port monitor type. You may not specify both this option and the **-p** option. This option adds the service to all port monitors of type *type*.

-i *login* Specify the user login name that will own the port service process. The login name must already exist.

-m *options* Use this option to specify port monitor-specific options for **pmadm**. To generate the options, embed a port monitor-specific command, delimited with shell backquote characters (`), as an argument to the **-m** option. The port monitor-specific command for **ttymon** is **ttyadm**(1M). The port monitor-specific command for **listen** is **nlsadmin**(1M).

-v *version* The version of the port monitor administrative file. For a port monitor of type **listen**, for example, the version number may be given as

```
-v `nlsadmin -V`
```

The version stamp of the port monitor is known by the port monitor-specific command and is returned when the command is invoked with the **-V** option.

-f Specify one or both of two flags which are then included in the flags field of the port monitor administrative file entry for the new service. The flags have these meanings:

x Do not enable the service.

u Create a **utmp** entry for the service.

If the **-f** option is not included on the **-a** command line, no flags are set and the default conditions prevail. By default, a new service is enabled and no **utmp** entry is created for it.

-y *"comment"* Specify a comment in double quotes. The comment is included in the comment field for the service entry in the port monitor administrative file.

-z *script* Installs *script* as a configuration file.

The following example adds a service with service tag **105** to all port monitors of type **listen**. The line is broken for readability.

```
# pmadm -a -s 105 -i root -t listen -v `nlsadmin -V` \
          -m `nlsadmin -a 105 -c /usr/net/servers/rfsetup` ⤶
```

Enabling or Disabling a Service

```
pmadm -e -p pmtag -s svctag
pmadm -d -p pmtag -s svctag
```

The **–e** option enables a service. **x** is removed from the flags field in the entry for service *svctag* in the port monitor administrative file.

The **–d** option disables a service. **x** is added to the flags field in the entry for service *svctag* in the port monitor administrative file.

Removing a Service

```
pmadm -r -p pmtag -s svctag
```

The **–r** removes service *svctag*. The entry for the service is removed from the port monitor administrative file.

Printing, Installing, and Replacing Per-Service Configuration Scripts

```
pmadm -g -p pmtag -s svctag [ -z script ]
pmadm -g -s svctag -t type -z script
```

Per-service configuration scripts are command scripts written in the interpreted language described in the **doconfig**(3N) manual page. They allow the system administrator to modify the environment in which a service executes. For example, the values of environment variables may be changed, STREAMS modules may be specified, or commands may be run.

Per-service configuration scripts are interpreted by the port monitor before the service is invoked. SAC interprets both its own configuration file, **_sysconfig**, and the port monitor configuration files. Only the per-service configuration files are interpreted by the port monitors. Per-service configuration scripts may be printed by any user on the system. Only the system administrator may install or replace them.

The **–g** option is used to print, install, or replace a per-service configuration script. The **–g** option with a **–p** option and a **–s** option prints the per-service configuration script for service *svctag* available through port monitor *pmtag*. The **–g** option with a **–p** option, a **–s** option, and a **–z** option installs the per-service configuration script contained in the file *script* as the per-service configuration script for service *svctag* available through port monitor *pmtag*. The **–g** option with the **–s** option, a **–t** option, and a **–z** option installs the file *script* as the per–service configuration script for service *svctag* available through any port monitor of type *type*. Other combinations of options with **–g** are invalid.

Examine the following sample per-service configuration script:

```
runwait ulimit 4096
runwait umask 077
```

This script does two things. It specifies the maximum file size for files created by a process by setting the process's **ulimit** to **4096**. It also specifies the protection mask to be applied to files created by the process by setting **umask** to **077.**

Appendix C contains a reference table that you may find useful for managing port services.

The Port Monitor ttymon

ttymon is a port monitor invoked by the Service Access Controller (SAC). The Service Access Controller is the Service Access Facility's controlling process. It is started by **init** when the system enters multiuser mode. One of SAC's functions after it is started is to start all port monitors the system administrator has configured.

ttymon sets terminal modes and line speeds for the port the user is connected to, allowing communication with the service associated with that port.

What ttymon Does

ttymon has three main functions:

- It initializes and monitors TTY ports.
- It sets terminal modes and line speeds for each port it monitors.
- It invokes the service associated with a given port whenever it receives a connection request on that port.

Each instance of **ttymon** has its own administrative file that specifies the ports to monitor and the services associated with each port. The file contains a *ttylabel* field that refers to a speed and TTY definition in the **/etc/ttydefs** file. See **ttyadm** (1M) for a description of the information specific to **ttymon** that is contained in a **ttymon** administrative file.

When a **ttymon** port monitor is started, it initializes all ports specified in its administrative file, pushes the specified STREAMS modules onto the ports, sets speed and initial **termio**(7) settings, and writes the prompt to the port. It then waits for user input.

A connection request is successful when at least one non-break character followed by a New Line character is received from the port. If the service to be invoked is **login**, the New Line character will be preceded by the users login name. A New Line character will not be recognized unless the line speed of the port and the line speed of the device connected to the port are the same.

If an unreadable prompt is printed on the terminal, the user sends a BREAK to indicate that the port and device line speeds are not compatible. See your terminal documentation to see how to generate a BREAK signal. Each break indication will cause **ttymon** to hunt to the next *ttylabel* in **/etc/ttydefs**, adjusting its **termio**(7) values and reissuing the prompt.

On successful completion of the connection request, **ttymon** interprets the per-service configuration script, if one exists. It then invokes the service associated with the port. This service can be any service configured by the system administrator. A typical example is **login**.

ttymon has no interaction with its TTY ports while they are connected to a service. On completion of a service on a port, **ttymon** returns the port to its initial state.

The Autobaud Option

Autobaud allows the system to set the line speed of a given TTY port to the line speed of the device connected to the port without the user's intervention. Each time a service to be monitored by a **ttymon** port monitor is added, a *ttylabel* must be supplied (see "Adding a Service," below). If this *ttylabel* points to an entry in the **/etc/ttydefs** file that has an **A** in the autobaud field, **ttymon** will try to determine the proper line speed before printing the prompt.

After receiving a carrier-indication on one of its TTY ports, but before printing a prompt, **ttymon** does the following:

- **ttymon** reads the next character received from the port. Provided the character read is a New Line character and that it is transmitted at a line speed autobaud can support, **ttymon** will reliably determine this line speed and change the port's line speed to that speed.

- If a baud rate cannot be determined from the character that is read (for example, if the user entered a character other than a New Line), or if a break is received rather than a character, **ttymon** considers this to be an autobaud failure and the character is discarded. If after five opportunities, a New Line is not recognized, the search proceeds to the next **ttydefs** entry in the hunt sequence. If an autobaud flag is encountered again, the prompt will not be written and the procedure just described is repeated. If no autobaud flag is set, the search again proceeds to the next **ttydefs** entry in the hunt sequence.

ttymon and the Service Access Facility

The Service Access Facility (SAF) provides a generic interface to which all port monitors must conform. **ttymon** is a port monitor under the Service Access Facility's controller, the Service Access Controller. (See "Overview of the Service Access Facility," "Port Monitor Management," and "Service Management" for a description of the Service Access Facility, the administrative files it maintains, and the commands used for port monitor and service administration.) Figure 9–1 shows how a service, which may be a **login** service, is invoked using **ttymon**.

There can be multiple invocations of **ttymon** port monitors, each identified by a unique *pmtag*. Each of these port monitors can monitor multiple ports for incoming connection requests.

A port has one and only one service associated with it. Each port and its associated service are identified by a service tag, *svctag*. Service tags for any given port monitor are unique.

When the Service Access Controller starts a port monitor, the port monitor reads its administrative file, which contains information about which ports to monitor and what service (that is, process) is associated with each port.

The Default ttymon Configuration

Some **ttymon** port monitors may be set up automatically when the system goes to multiuser level. To find out if your system has been automatically configured, enter the command

```
# sacadm -l ↵
```

after the system is in multiuser mode. To see a listing of all services available under the configured **ttymon** port monitors, enter the command

```
# pmadm -l -t ttymon ↵
```

The line discipline module, **ldterm**, may not be specified for automatically configured services. Instead, it may be defined in an autopush administrative file and pushed by the autopush facility (see the **autopush**(1M) manual page). **autopush** pushes previously specified modules onto the appropriate STREAM each time a device is opened.

Services are not defined for the console and contty ports under any **ttymon** port monitor. Instead, there is an entry for each in the **/etc/inittab** file. These entries contain calls to **ttymon** in "express." See "**ttymon** Express," below.

The ttyadm Command

The Service Access Facility requires each type of port monitor to provide an administrative command. This command must format information derived from command-line options so that it is suitable for inclusion in the administrative files for that port monitor type. The command may also perform other port monitor-specific functions.

ttyadm is **ttymon** 's administrative command. The **ttyadm** command formats information based on the options with which it is invoked and writes this information to the standard output.

ttyadm is one of the arguments **pmadm** uses with the **–a** option to format information in a way suitable for inclusion in a **ttymon** administrative file. **ttyadm** presents this information (as standard output) to **pmadm**, which places it in the file. This use of **ttyadm** is described below under "Adding a **ttymon** Port Monitor." Port monitor-specific information in a port monitor administrative file will be different for different port monitor types.

ttyadm is also included on the **sacadm** command line when a port monitor is added to the system. It is used to supply the **ttymon** version number for inclusion in a port monitor's administrative file. The port monitor administrative file is updated by the Service Access Controller's administrative commands, **sacadm** and **pmadm**. **ttyadm** merely provides a means of presenting formatted port monitor-specific (that is, **ttymon**-specific) data to these commands. The **sacadm** command line uses **ttyadm** only with the **–V** option. **ttyadm –V** tells SAC the version number of the **ttymon** command being used.

Managing TTY Ports

Listing Configured ttymon Port Monitors

```
sacadm -l  [ -p pmtag | -t type ]
```

The **sacadm** command with only a **-l** option lists all port monitors currently defined for the system. The following is an example of its output:

```
PMTAG     PMTYPE     FLGS RCNT STATUS     COMMAND
tcp       listen     -    3    ENABLED    /usr/lib/saf/listen tcp #listener
for tcp
ttymon1   ttymon     -    0    ENABLED    /usr/lib/saf/ttymon #
```

sacadm can also be used to list a single port monitor (**-p**) or to list only port monitors of a single type (**-t**), for example, all port monitors of type **ttymon**. For a complete description of these options, see "Printing Port Monitor Status Information" in "Port Monitor Management" above, or see the **sacadm**(1M) manual page.

Listing Services Configured for a ttymon Port Monitor

```
pmadm -l  [-p pmtag | -t type] [ -s svctag]
```

pmadm with only a **-l** will list all services for all port monitors on the system. If a port monitor is specified (**-p**), all services for that port monitor will be listed. The following is a sample listing for the command

```
# pmadm -l -p ttymon2 ↵
```

```
PMTAG    PMTYPE SVCTAG FLGS ID   PMSPECIFIC
ttymon2 ttymon 21     ux   root /dev/tty21 -- /usr/bin/login - 9600 -
login: -
ttymon2 ttymon 22     ux   root /dev/tty22 -- /usr/bin/login - 9600 -
login: -
ttymon2 ttymon 23     ux   root /dev/tty23 -- /usr/bin/login - 9600 -
login: -
ttymon2 ttymon 24     ux   root /dev/tty24 -- /usr/bin/login - 9600 -
login: -
```

In the above table, the *pmspecific* fields include the device (for example, **/dev/tty21**), the service to be invoked (**/usr/bin/login**), and the prompt (`login:`). See the **ttyadm**(1M) manual page for a description of the *pmspecific* fields.

Listing Accessible TTY Ports

To find out which ports are accessible to users, first identify all enabled **ttymon** port monitors:

```
# sacadm -l -t ttymon ↵
PMTAG     PMTYPE     FLGS RCNT STATUS     COMMAND
ttymon1   ttymon     -    0    ENABLED    /usr/lib/saf/ttymon #
```

In the listing, port monitor **ttymon1** is enabled. This means that it is accepting service requests for any of its services that are enabled.

To identify which services are enabled, use **pmadm –l –p ttymon1**. This will list all configured TTY services for port monitor **ttymon1** as in the following example:

```
# pmadm  -l  -p  ttymon1 ↵
PMTAG   PMTYPE SVCTAG FLGS ID  <PMSPECIFIC>
ttymon1 ttymon 11     u    root /dev/tty11 - - /usr/bin/login - 9600 -
login: -
ttymon1 ttymon 12     ux   root /dev/tty12 - - /usr/bin/login - 9600 -
login: -
ttymon1 ttymon 13     u    root /dev/tty13 - - /usr/bin/login - 9600 -
login: -
ttymon1 ttymon 14     ux   root /dev/tty14 - - /usr/bin/login - 9600 -
login: -
```

In the listing, enabled services are those that do *not* have an **x** in the FLGS column. The ports corresponding to these services (**/dev/tty11** and **/dev/tty13**) are accessible to users. The **who –l** command lists all running port monitors, not the accessible TTY ports. Follow the procedure described above to find out which TTY ports are accessible.

Adding a ttymon Port Monitor

```
sacadm -a -p pmtag -t type -c cmd -v `ttyadm -v` \
                  -n count [ -f dx ] [ -z script ] [ -y comment ]
```

The following command line will add a **ttymon** -type port monitor named **ttymon1**:

```
# sacadm -a -p ttymon1 -t ttymon -c /usr/lib/saf/ttymon \
       -v `ttyadm -v` ↵
```

The command adds a line to SAC's administrative file. The options that may be used with **sacadm –a** are described under "Port Monitor Management," above, and in the **sacadm**(1M) and **ttyadm**(1M) manual pages. If a port monitor already exists with the same name as the port monitor that is being added, the system administrator must remove the old port monitor before adding the new one.

Removing a ttymon Port Monitor

```
sacadm -r -p pmtag
```

The following command line removes the port monitor added in the previous example:

```
# sacadm -r -p ttymon1 ↵
```

SAC removes the line for port monitor **ttymon1** from its administrative file. The port monitor directory will remain in **/etc/saf** but will be removed and recreated when a new port monitor with the same name is added. To make changes to a port monitor entry, always remove the entry and add a new entry using **sacadm** or **sysadm**.

NOTE: Do not edit the SAC administrative file. If you edit the file, you could introduce format or syntax errors that could affect the function of your port monitors in undesirable ways.

Adding a Service

```
pmadm -a -p pmtag -s svctag -i id [ -f ux ] -v `ttyadm -V` -m "`ttyadm \
    [ -b ] [ -r count ] [ -c ] [ -h ] [ -i msg ] [ -m modules ] \
    [ -p prompt ] [ -t time-out ] -d device -l ttylabel -s service `"
```

The following command line adds a login service to be monitored by the **ttymon** port monitor **ttymon2**:

```
# pmadm -a -p ttymon2 -s 21 -i root -fu -v `ttyadm -V` -m"`ttyadm -d \
    /dev/tty21 -l 9600 -s /usr/bin/login -m ldterm -p \"tty21:\"`" ↵
```

The options that may be used with **pmadm –a** are described under "Service Management," above, and on the **pmadm**(1M) and **ttyadm**(1M) manual pages.

The **ttyadm –m** option may be used for pushing STREAMS modules, for example the line discipline module, **ldterm**. If **autopush** has pushed modules on the stream, **ttymon** pops them before pushing its own.

By using the **ttyadm –i** option, we could also have specified a message to be printed whenever someone tries to log in on a disabled port.

The following command defines a service that permits both incoming and outgoing calls. The service is put under port monitor **ttymon2**. The **–b** option defines the port as bidirectional.

```
# pmadm -a -p ttymon2 -s 21 -i root -fu -v `ttyadm -V` \
     -m "`ttyadm -b -h -r0 -t 60 -d /dev/tty21 \
     -l 9600H -s /usr/bin/login -m ldterm -p \" tty21:\"`" ↵
```

The **ttyadm –r** option with **count=0** is assumed when the **ttyadm –b** bidirectional option is used; the **–r0** could therefore have been omitted.

Removing a Service

```
pmadm  -r  -p  pmtag -s  svctag
```

The following example deletes the service that was added in the previous example.

```
# pmadm -r -p ttymon2 -s 21 ↵
```

Enabling a Service

```
pmadm  -e  -p  pmtag  -s  svctag
```

To enable a service on a specific port, first find out which port monitor is monitoring the port. Enter

```
# pmadm -l -t ttymon ↵
```

This lists all services defined for ttymon-type ports.

Now look in the PMSPECIFIC column for the device file that corresponds to the port you are interested in, for example, **/dev/tty23**. If the port monitor is **ttymon2** and the service tag is **23**, the command

```
# pmadm -e -p ttymon2 -s 23 ↵
```

will enable the service on port **/dev/tty23**.

To verify that the port has been enabled, enter

```
# pmadm -l -p ttymon2 -s 23 ↵
```

The **x** will have been removed from the FLGS column in the entry for this service.

Disabling a Service

```
pmadm -d -p pmtag -s svctag
```

When a service is disabled, all subsequent connection requests for the service will be denied. Using the same example,

```
# pmadm -d -p ttymon2 -s 23 ↵
```

will restore the **x** to the FLGS field in the entry for service **23**.

Disabling All Services Monitored by a ttymon Port Monitor

```
sacadm -d -p pmtag
```

To disable all services defined for the port monitor **ttymon2**, enter

```
# sacadm -d -p ttymon2 ↵
```

Any future connection requests for services managed by this port monitor will be denied until the port monitor is enabled.

The command

```
# sacadm -e -p ttymon2 ↵
```

will re-enable port monitor **ttymon2**.

ttymon Express Mode

Services are not defined for the console and TTY ports under any **ttymon** port monitor. Instead, there is an entry for each in the **/etc/inittab** file. These entries contain calls to **ttymon** in "express" mode. **ttymon** express is a special mode of **ttymon** that permits **ttymon** to be invoked directly by a command that requires login service. **ttymon** in express mode is not managed by the Service Access Controller nor is an administrative file associated with any invocation of **ttymon** in this mode.

ttymon express is described in greater detail on the **ttymon**(1M) manual page.

Configuration Files

As a port monitor under the Service Access Facility, **ttymon** can customize the environment of each service it starts. It does this by interpreting a per-service

configuration script, if one exists, immediately before starting the service. Per-service configuration scripts are optional. Configuration scripts are installed by the system administrator, using the **pmadm** command with **–g** and **–z** options (see the **pmadm**(1M) manual page).

It is also possible to customize the environment of a **ttymon** port monitor. A per-port monitor configuration script is defined using the **sacadm** command with **–g** and **–z** options (see the **sacadm**(1M) manual page). The environment modifications made by a port-monitor configuration script are inherited by the port monitor and all the services it invokes. The environment of any particular service can then be customized further by using a per-service configuration script.

The **doconfig**(3N) manual page describes the language in which configuration scripts are written.

Configuration scripts are not normally needed for basic operations.

The who Command

The **who** command examines the **/etc/utmp** file. It is used to find out who is on the system. The following command lists all entries in the **utmp** file including all **RUNNING** port monitors:

```
# who -lH ↵

NAME     LINE     TIME          IDLE  PID   COMMENTS
LOGIN    contty   Jun 17 12:49  old   8226
tcp      .        Jun 17 12:50  old   8230
ttymon1  .        Jun 17 12:50  old   8234
ttymon3  .        Jun 17 12:50  old   8235
```

When **ttymon** in express mode is monitoring a line, the name field is **LOGIN** as it is in the entry for **contty** in the preceding example.

The following command lists all users who are currently logged in:

```
# who -u ↵
root       console      Jun 17 13:07   .      8303
john       term/32      Jun 17 13:13   0:01   8353
```

If **ttymon** invokes a service other than **login**, an entry for this service will appear beginning with a "**\.**" and giving the terminal line. To find out which ports are accessible but currently not in use, use the following command:

```
# pmadm -l -t ttymon ↵
```

Identifying ttymon Processes

The **ps** command lists all processes. Since **ttymon** port monitors fork a process to handle each connection request, the number of **ttymon**-related entries that appear in the output of a **ps** listing may be greater than the number of running **ttymon** port monitors.

When a **ttymon** port monitor forks a child to process a connection request (that is, to do baud rate searching, set final **termio** options, and so on, before invoking the service), the port will be identified in the TTY field for this child process. For the parent **ttymon** port monitor process, this TTY field will be empty.

Log Files

Problems often arise when a single port is monitored by more than one process. If a port (for example, **/dev/tty11**) is used by an enabled service under a **ttymon** port monitor running under the Service Access Facility, and the same port is also monitored by a **ttymon** process running in **ttymon** express mode, (that is, started by **init** when it reads **inittab**, not by **sac** when it reads its administrative file) then the port will behave unpredictably. The system administrator is expected to examine the system for such ambiguously configured ports.

There are also two log files that can be examined for clues to problems related to **ttymon** port monitors or ports monitored by **ttymon** port monitors: The Service Access Controller records aberrant port monitor behavior in **/var/saf/_log**; and each **ttymon** port monitor has its own log file, **/var/saf/***pmtag***/log**, where it records messages it receives from SAC, services it starts, and so on.

The following command:

```
# tail  -25  /var/saf/_log ↵
```

will list the most recent 25 entries in the **_log** file.

Periodically, you should truncate the log files. You can set up a **cron** job to clean up regularly. Refer to Chapter 2 for information to automate jobs with **cron** and to perform log cleanup. Also see the manual pages for **cron**(1M) and **crontab**(1).

Terminal Line Settings

init is a general process spawner that is invoked as the last step in the boot procedure. It starts SAC. SAC then looks in its administrative file to see which port monitors to start. Each **ttymon** port monitor started by SAC looks in *its* administrative file for the TTY ports to initialize. For each TTY port initialized, **ttymon** searches the **ttydefs** file for the information it needs to set terminal modes and line speeds. **ttymon** then waits for service requests. When a service request is received, **ttymon** executes the command (usually **login**) associated with the port that received the request. This command is contained in the entry for the port in the port monitor's administrative file.

From the system administrator's point of view, the key elements in managing terminal line settings are the **ttydefs** file and the **sttydefs** command, which maintains the **ttydefs** file.

The sttydefs Command

sttydefs(1M) is an administrative command that maintains the **ttydefs** file. The **ttydefs** file contains information about line settings and hunt sequences for the

system's TTY ports. The **sttydefs** command and the **ttydefs** file together provide the facilities for managing terminal modes and line settings. The **sttydefs** command is used to:

- Print information contained in **ttydefs**.
- Add records for terminal ports to the **ttydefs** file.
- Remove records from the **ttydefs** file.

Printing Terminal Line Setting Information

```
/usr/sbin/sttydefs -l [ttylabel]
```

If a *ttylabel* is specified, **sttydefs** prints the **ttydefs** record that corresponds to this *ttylabel*. If no *ttylabel* is specified, **sttydefs** prints this information for all records in the **/etc/ttydefs** file. **sttydefs** verifies that each entry it displays is correct and that the entry's *nextlabel* field refers to an existing *ttylabel*. An error message is printed for each invalid entry detected.

Adding Records to the ttydefs File

```
/usr/sbin/sttydefs -a ttylabel [-b] [-n nextlabel] [-i initial-flags] \
        [-f final-flags]
```

sttydefs with the **–a** option adds a record to the **ttydefs** file.

ttylabel identifies the record.

The following list describes the effect of the **–b**, **–n**, **–i**, or **–f** options when used with the **–a** option. The **–a** option is valid only when invoked by a privileged user.

–b Enables autobaud.

–n Specifies the value to be used in the *nextlabel* field. If *nextlabel* is not specified, **sttydefs** will set *nextlabel* to *ttylabel*.

–i Specifies the value to be used in the *initial-flags* field. The argument to this option must be presented in a format recognized by the **stty** command. If *initial-flags* is not specified, **sttydefs** will set *initial-flags* to the **termio**(7) flag **9600**.

–f Specifies the value to be used in the *final-flags* field. The argument to this option must be presented in a format recognized by the **stty** command. If *final-flags* is not specified, **sttydefs** will set *final-flags* to the **termio**(7) flags **9600** and **sane**.

The following command line creates a new record in **ttydefs**:

```
# sttydefs -aNEW -nNEXT -i"1200 hupcl erase ^h" -f"1200 sane ixany \
        hupcl erase ^h echoe" ↵
```

The flag fields shown have the following meanings:

300–19200 The baud rate of the line.

hupcl Hang up on close.

sane	A composite flag that stands for a set of normal line characteristics.
ixany	Allow any character to restart output. If this flag is not specified, only DC1 (Ctrl-Q) will restart output.
tab3	Send tabs to the terminal as spaces.
erase ^h	Set the erase character to **^h**. On most terminals a **^h** is the backspace.
echoe	Echo erase character as the string backspace-space-backspace. On most terminals this will erase the erased character.

Creating a Hunt Sequence

The following sequence of commands adds records with labels **1200**, **2400**, **4800**, and **9600** to the **ttydefs** file and puts them in a circular list or hunt sequence:

```
# sttydefs -a1200 -n2400 -i 1200 -f "1200 sane" ↵
# sttydefs -a2400 -n4800 -i 2400 -f "2400 sane" ↵
# sttydefs -a4800 -n9600 -i 4800 -f "4800 sane" ↵
# sttydefs -a9600 -n1200 -i 9600 -f "9600 sane" ↵
```

The *nextlabel* field of each line is the *ttylabel* of the next line. The *nextlabel* field for the last line shown points back to the first line in the sequence.

The object of a hunt sequence is to link a range of line speeds. Entering a BREAK during the baud rate search causes **ttymon** to step to the next entry in the sequence. See your terminal documentation to see how to generate a BREAK signal. The hunt continues until the baud rate of the line matches the speed of the user's terminal.

The **ttydefs** file containing these records will look like this:

```
# VERSION=1
1200:1200:1200 sane::2400
2400:2400:2400 sane::4800
4800:4800:4800 sane::9600
9600:9600:9600 sane::1200
```

Removing Records from the ttydefs File

/usr/sbin/sttydefs -r *ttylabel*

The record for the *ttylabel* specified on the command line is removed from the **ttydefs** file.

The **–r** option is valid only when invoked by a privileged user. If a record you remove is part of a hunt sequence, be sure the sequence is repaired. It may be useful to run **sttydefs** with the **–l** option after a record has been removed. **sttydefs –l** will check for incorrect field values and broken hunt sequences and will print error messages.

The ttydefs File

/etc/ttydefs is an administrative file used by **ttymon**. It defines speed and terminal settings for TTY ports. The **ttydefs** file contains the information listed below. The figure following shows the relationship between the *ttylabel* and *nextlabel* fields in the **ttymon** administrative files and **ttydefs** files. The example after that shows a sample **ttydefs** file.

ttylabel When **ttymon** initializes a port, it searches the **ttydefs** file for the entry that contains the **termio**(7) settings for that port. The correct entry is the one whose *ttylabel* matches the *ttylabel* for the port. The *ttylabel* for the port is part of the *pmspecific* information included in **ttymon**'s administrative file. By convention, **ttylabel** identifies a baud rate (for example, **1200**), but it need not.

initial–flags Contains the **termio**(7) options to which the terminal is initially set. *initial–flags* must be specified using the syntax recognized by the **stty**(1) command.

final–flags Contains the **termio**(7) options set by **ttymon** after a connection request has been made and immediately before invoking a port's service. *final–flags* must be specified using the syntax recognized by **stty**.

autobaud *autobaud* is a line-speed option. When *autobaud* is used instead of a baud rate setting, **ttymon** determines the line speed of the TTY port by analyzing the first carriage return entered and sets the speed accordingly. If the *autobaud* field contains the character **A**, the *autobaud* facility is enabled; otherwise, *autobaud* is disabled.

nextlabel If the user indicates (by sending a BREAK) that the current **ttydefs** entry does not provide a compatible line speed, **ttymon** will search for the **ttydefs** entry whose *ttylabel* matches the *nextlabel* field. **ttymon** will then use that field as its *ttylabel* field. A series of speeds is often linked together in this way into a closed set called a hunt sequence. For example, **4800** may be linked to **1200**, which in turn is linked to **2400**, which is finally linked to **4800**.

All **termio**(7) settings supported by the **stty** command are supported as options in the **ttydefs** file. For example, the system administrator will be able to specify the default erase and kill characters.

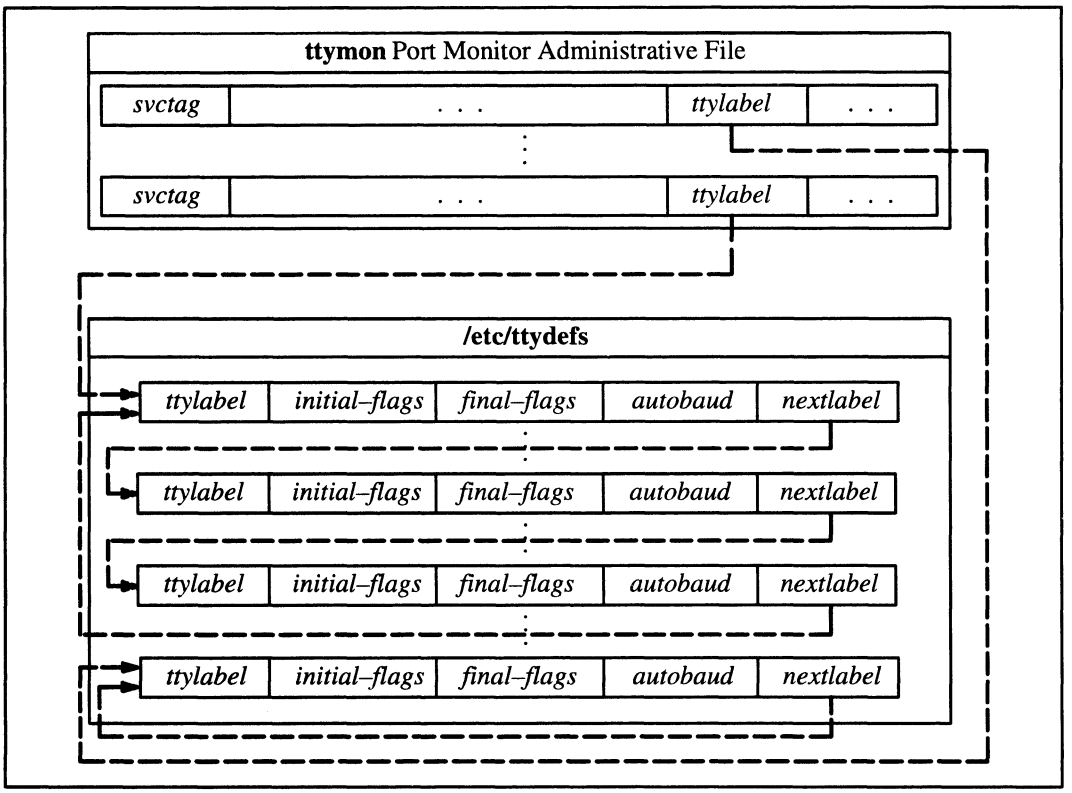

Figure 10–2 Port monitor/ttydefs Links

The format of the **/etc/ttydefs** file may change in future releases. For continuity across releases, use the **sttydefs**(1M) command to access this file. Following is a sample **ttydefs** file:

```
# VERSION=1
38400:38400 hupcl erase ^h:38400 sane ixany tab3 hupcl erase ^h::19200
19200:19200 hupcl erase ^h:19200 sane ixany tab3 hupcl erase ^h::9600
9600:9600 hupcl erase ^h:9600 sane ixany tab3 hupcl erase ^h::4800
4800:4800 hupcl erase ^h:4800 sane ixany tab3 hupcl erase ^h::2400
2400:2400 hupcl erase ^h:2400 sane ixany tab3 hupcl erase ^h::1200
1200:1200 hupcl erase ^h:1200 sane ixany tab3 hupcl erase ^h::300
300:300 hupcl erase ^h:300 sane ixany tab3 hupcl erase ^h::19200
```

Setting Terminal Options with the stty Command

The **stty**(1) command may be used to set or change terminal options after a user has logged in. A **stty** command line may also be added to a user's **.profile** to set options automatically as part of the **login** process. The following is an example of a simple **stty** command:

```
# stty cr0 nl0 echoe -tabs erase ^h ↵
```

The options in the example have the following meanings:

cr0 nl0 No delay for carriage return or new line. Delays are not used on a video display terminal, but are necessary on some printing terminals to allow time for the mechanical parts of the equipment to move.

echoe Erase characters as you backspace.

–tabs Expand tabs to spaces when printing.

erase ^h Change the character-delete character to **^h**. The default character-delete character is the pound sign (**#**). Most terminals transmit a **^h** when the **backspace** key is pressed.

Appendix C contains a reference table that you may find useful for setting terminal lines.

The Port Monitor listen

listen is a port monitor invoked by the Service Access Controller (SAC). The Service Access Controller is the Service Access Facility's controlling process. It is started by **init** when the system enters multiuser mode. One of SAC's functions after it is started is to start all port monitors the system administrator has configured.

listen monitors a connection-oriented transport network, receiving incoming connection requests, accepting them, and invoking the services that have been requested. The **listen** port monitor may be used with any connection-oriented transport provider that conforms to the Transport Interface (TLI) specification. The TLI is documented in *Programming with TCP/IP on the DG/UX™ System.*

What listen Does

The **listen** port monitor performs functions common to all port monitors:

- It initializes and monitors **listen** ports, and
- it invokes the service associated with a port in response to requests.

The **listen** port monitor also has these features:

- It allows private addresses for services,
- passes connections (file descriptors) to standing servers,
- supports socket-based services, and
- supports RPC-based services and dynamic addressing.

Private Addresses for Services

Each **listen** service may have a transport address in addition to its service code (*svctag*). This private address is included in the port monitor's administrative file. The inclusion of private addresses for services allows a single **listen** process to monitor multiple addresses. The number of addresses that **listen** can listen on is determined by the number of file descriptors available to the process.

Passing Connections to Standing Servers

By default, a new instance of a service is invoked for each connection. This feature is useful for server processes that need to maintain state information. A standing server is a server process or service that runs continuously and accepts connections through a FIFO or a named STREAM instead of being propagating copies of itself with calls to **fork**(2) and **exec**(2).

Socket-based Services

listen supports services that use sockets as their interface to the transport provider. Socket-based services are registered with **listen** in the same way TLI services are, using the Service Access Facility's administrative commands. **listen** supports STREAMS; sockets are implemented as a STREAMS module and a library.

A socket-based service:

- May or may not be an RPC service.
- May have a statically or dynamically assigned address, or no private address.
- May be invoked on each connection or may be a standing server, to which new connections are passed by a FIFO or a named STREAM.

RPC Services and Dynamic Addressing

Dynamic addressing is most useful with RPC. RPC transport addresses may be either specified or dynamically assigned. In either case, **listen** tells the **rpcbinder** what the address is and monitors it for incoming connections.

In the case of a dynamically assigned address, **listen** asks the transport provider to select a transport address each time **listen** begins listening on behalf of the service.

When service addresses are dynamically assigned, the assigned address is written to the **listen** log file.

listen and the Service Access Facility

The Service Access Facility (SAF) provides a generic interface to which all port monitors must conform. **listen** is a port monitor under the Service Access Facility's controller, the Service Access Controller. (See "Overview of the Service Access Facility," "Port Monitor Management," and "Service Management," above, for a description of the Service Access Facility, the administrative files it maintains, and the commands used for port monitor and service administration.)

There can be multiple invocations of **listen** port monitors, each identified by a unique *pmtag*. Each of these port monitors can monitor multiple ports for incoming connection requests.

A port has one and only one service associated with it. Each port, and its associated service, is identified by a service tag, *svctag*. Service tags for any given port monitor are unique.

When the Service Access Controller starts a port monitor, the port monitor reads its administrative file, which contains information about which ports to monitor and what service (that is, process) is associated with each port.

The nlsadmin Command

The Service Access Facility requires each type of port monitor to provide an administrative command. This command must format information derived from command-line options so that it is suitable for inclusion in the administrative files for that port monitor type. The command may also perform other port monitor functions.

nlsadmin is **listen**'s administrative command. The **nlsadmin** command formats information based on the options with which it is invoked and writes this information to the standard output.

nlsadmin is one of the arguments **pmadm** uses with the **–a** option to format information in a way suitable for inclusion in a **listen** administrative file. **nlsadmin** presents this information (as standard output) to **pmadm**, which places it in the file. This use of **nlsadmin** is described below under "Adding a **listen** Port Monitor." Port monitor specific information in a port monitor administrative file will be different for different port monitor types.

nlsadmin is also included on the **sacadm** command line when a port monitor is added to the system. It is used to supply the **listen** version number for inclusion in a port monitor's administrative file. The port monitor administrative file is updated by the Service Access Controller's administrative commands, **sacadm** and **pmadm**. **nlsadmin** merely provides a means of presenting formatted port monitor-specific data to these commands.

The **sacadm** command line uses **nlsadmin** only with the **–V** option. **nlsadmin –V** tells SAC the version number of the **listen** command being used.

Under the SAF, it is possible to have multiple instances of **listen** on a single *net_spec*. A new option, **–N** *pmtag*, can be used in place of the *net_spec* argument. This argument specifies the tag by which an instance of **listen** is identified by the SAF. If the **–N** option is not specified (i.e. the *net_spec* is specified in the invocation), then it will be assumed that the last component of the *net_spec* represents the tag of **listen** for which the operation is destined.

Managing listen Ports

Listing Configured listen Port Monitors

```
# sacadm  -l  [ -t listen ]
```

The **sacadm** command with only a **–l** option lists all port monitors currently defined for the system. For example:

```
PMTAG    PMTYPE  FLGS RCNT STATUS   COMMAND
tcp      listen  -    3    ENABLED /usr/lib/saf/listen tcp #listener for tcp
ttymon1 ttymon  -    0     ENABLED /usr/lib/saf/ttymon #
```

Listing Services Configured for a listen Port Monitor

```
pmadm -l [-p net_spec] [-s svctag]
```

pmadm with only a **–l** will list all services for all port monitors on the system. If a port monitor is specified (**–p**), all services for that port monitor will be listed. The following is a sample listing for the command.

```
# pmadm  -l  -p  tcp ↵

PMTAG PMTYPE SVCTAG FLGS ID     PMSPECIFIC
tcp   listen 101    -    listen - c - /usr/lib/uucp/uucico -r0 -unuucp -iTLI \
                                      #UUCP access direct to server
tcp   listen 102    -    listen - c - /usr/tcp/lib/ttysrv -d -n ntty,tirdwr,1d0 \
                                      #UUCP access to server via login
tcp   listen 1000   -    listen - c - /usr/tcp/lib/tstserver #TP TEST SERVER
tcp   listen 0      -    root   - c - sfbul.serve c - /usr/lib/saf/nlps_server \
                                      #NLPS server
```

The following command lines list the addresses associated with general **listen** service (**0**) or with login service (**1**).

```
pmadm -l -p net_spec -s 0
pmadm -l -p net_spec -s 1
```

By definition, service code **0** is for the **nlps_server**, which is a service that provides compatibility with pre-DG/UX Release 5.4 **listen** service requests. Service code **1** is for remote login (that is, **cu** over a network).

Adding a listen Port Monitor

```
sacadm -a -p pmtag -t type -c cmd -v `pmspecific -v` \
            -n count [ -f dx ] [ -z script ] [ -y comment ]
```

The following example shows how **listen**'s administrative command, **nlsadmin**, can be used to obtain the current version number of **listen**'s administrative file when used with **sacadm** to add a **listen** port monitor.

```
# sacadm -a -p tcp -t listen -c "/usr/lib/saf/listen -m tcp" \
         -v `nlsadmin -V` ↵
```

This command line adds a line to SAC's administrative file. The options that may be used with **sacadm –a** are described under "Port Monitor Management," above, and in the **sacadm**(1M) and **nlsadmin**(1M) manual pages. If the port monitor being added has the same name as an existing port monitor, the system administrator must remove the old one before adding the new one.

Removing a listen Port Monitor

```
sacadm -r -p net_spec
```

For example,

```
# sacadm -r -p tcp ↵
```

SAC removes the line for port monitor **tcp** from its administrative file. The port monitor directory will remain in **/etc/saf** but will be removed and recreated when a new port monitor with the same name is added. To make changes to a port monitor entry, always remove the entry and add a new entry using the **sacadm** command. You are advised against editing the SAC administrative file.

Adding a Service

```
pmadm -a -p { net_spec | pmtag } -s svctag -i id -m "`nlsadmin options `" \
      -v `nlsadmin -V` -y comment
```

The following example adds a new service, **/usr/bin/cmd**, to a **listen** port monitor whose tag is **listen**. The new service has service tag **23**, identity **guest**, and no private address:

```
# pmadm -a -p listen -s 23 -i guest -m `/usr/sbin/nlsadmin \
     -c /usr/bin/cmd` -v `/usr/sbin/nlsadmin -V` ↵
```

The same address cannot be monitored by more than one **listen** port monitor at any given time. The first attempt to listen on an address will bind successfully; subsequent attempts will fail to bind. If both static and dynamic addresses are monitored by more than one **listen** port monitor, the static addresses are bound first, then the dynamic addresses. Mixing multiple **listen** port monitors—each of which has static and dynamic addresses specified—may result in unpredictable behavior. See "Adding a Service" under "Service Management," or the **pmadm**(1M) and **nlsadmin**(1M) manual pages for a full description of the **pmadm** command line options.

Removing a Service

```
pmadm -r -p { net_spec | pmtag } -s svctag
```

For example,

```
# pmadm -r -p tcp -s 23 ↵
```

removes service **23** from the **tcp listen** port monitor.

Enabling and Disabling Services

```
pmadm -e -p net_spec -s svctag
pmadm -d -p net_spec -s svctag
```

To enable a service on a specific port, first find out which port monitor is monitoring the port. Enter

```
# pmadm -l -t listen ↵
```

This lists all services defined for listen-type ports.

If the port monitor is **tcp** and the service tag is **101**, the command

```
# pmadm -e -p tcp -s 101 ↵
```

will enable service **101** . To verify that the port has been enabled, enter

```
# pmadm -l -p tcp -s 101 ↵
```

The **x** will have been removed from the FLGS column in the entry for this service. When a service is disabled, all subsequent connection requests for the service will be denied. Using the same example,

```
# pmadm -d -p tcp -s 101 ↵
```

will restore the **x** to the FLGS field in the entry for service 101.

Disabling All Services Monitored by a listen Port Monitor

sacadm -d -p *pmtag*

To disable all services defined for the port monitor **tcp**, enter

```
# sacadm -d -p tcp ↵
```

Any future connection requests for services managed by this port monitor will be denied until the port monitor is enabled. The command

```
# sacadm -e -p tcp ↵
```

will re-enable port monitor **tcp**.

Configuration Files

As a port monitor under the Service Access Facility, **listen** can customize the environment of each service it starts. It does this by interpreting a per-service configuration script, if one exists, immediately before starting the service. Per-service configuration scripts are optional. Configuration scripts are installed by the system administrator, using the **pmadm** command with **–g** and **–z** options (see the **pmadm**(1M) manual page).

It is also possible to customize the environment of a **listen** port monitor. A per-port monitor configuration script is defined using the **sacadm** command with **–g** and **–z** options (see the **sacadm**(1M) manual page). The environment modifications made by a port-monitor configuration script are inherited by the port monitor and all the services it invokes. The environment of any particular service can then be customized further by using a per-service configuration script.

The **doconfig**(3N) manual page describes the language in which configuration scripts are written.

Configuration scripts are not normally needed for basic operations.

Log Files

listen creates and manages the log files **/var/saf/***pmtag***/log** and **/var/saf/***pmtag***/o.log**. Log file entries are in the following format:

date time; PID; message

where
date and *time* show when the entry was made.

PID is the ID of the process that made the log entry.
message gives a description of the event or error that caused the log message.

The following events are logged:

- Each connection that arrives
- Each service that is started
- Each file descriptor passed over a pipe
- State changes that occur
- Errors and unusual conditions

The log files are held open by **listen** processes. Entries are made by two types of processes: **listen** process (**listen**) and the NLPS server process (**nlps_server**). **nlps_server** is a service that provides compatibility with pre-DG/UX Release 5.4 service requests.

Appendix C contains a reference table that you may find useful for managing **listen** port monitors.

End of Chapter

Chapter 11
Controller Management

This chapter tells how to manage the following kinds of controllers:

- Synchronous VSC controllers, used for wide-area networks (WANs).
- VLC controllers, used for local-area networks (LANs) .
- Asynchronous VDA controllers, used for terminal lines and other asynchronous connections.

Sync Management

This section describes how to manage your system's synchronous wide-area network (WAN) controllers. The supported sync controllers include the VSC/3, VSC/3i, and VSC/4 (VME Bus Synchronous Controller) controllers. The **sysadm** Device –> Sync menu contains operations for starting, stopping, checking, and listing your sync controllers.

You can use the Start, Stop, Check, and List operations only if intelligent synchronous communications drivers are configured on your system. If the controllers are installed on the system when you build your kernel, the system file will include the correct driver entries. The VSC/3 and VSC/4 controllers require the **ssid** driver, and the VSC/3i controller requires the **vsxb** driver.

The **sysadm** `Device -> Sync` operations are appropriate only for upper layer communications software such as X.25 and SNA. These operations are not appropriate for integrated synchronous controllers such as **iscd** and **izscd**.

Starting Sync Controllers

Use the **sysadm** operation `Device -> Sync -> Start` to download the controller-resident software to sync drivers and initialize them. The operation prompts you for the controllers that you wish to start. Select **all** to start all controllers.

You may not perform this operation on a controller if any port on the controller is in use.

Stopping Sync Controllers

Use the **sysadm** operation `Device -> Sync -> Stop` to halt all controller-resident software running on a sync controller. The operation stops the

controller by performing a hardware reset on the controller board. The operation prompts you for the sync controllers that you wish to stop. Select **all** to stop all controllers.

You may not perform this operation on a controller if any port on the controller is in use.

Checking Sync Controllers

Use the **sysadm** operation `Device -> Sync -> Check` to verify that sync controllers are functional. The operation prompts you for the controllers that you wish to check. Select **all** to check all controllers.

This operation checks controllers by calling the **synccheck**(1M) command. The **synccheck** command tests a controller by verifying that controller-resident software has been downloaded and that the controller can perform DMA operations across the VME bus.

Listing Sync Controllers

Use the **sysadm** operation `Device -> Sync -> List` to display the **/dev** entries for downloadable sync controllers configured on your system. The display may include entries such as **ssid** and **vsxb**.

LAN Management

This section describes how to manage local-area network (LAN) controllers that execute controller-resident software downloaded from the host and list all LAN controllers. The **sysadm** `Device -> LAN` menu contains operations for starting, stopping, and listing your LAN controllers.

You can use the start and stop operations only if intelligent LAN communications drivers are configured on your system. If the controllers are installed on the system when you build your kernel, the system file will include the correct driver entries. The VLCi controller is intelligent and requires the **cien** driver.

Starting LAN Controllers

Use the operation `Device -> LAN -> Start` to download controller-resident code to the LAN controller and initialize it. The operation prompts you for the LAN controllers that you wish to start. Select **all** to start all controllers.

Stopping LAN Controllers

Use the operation `Device -> LAN -> Stop` to halt all tasks running on a LAN controller. The operation stops the controller by performing a hardware reset on the controller board. The operation prompts you for the LAN controllers that you wish to stop. Select **all** to stop all controllers.

Listing LAN Controllers

Use the operation `Device -> LAN -> List` to display the **/dev** entries for your LAN controllers.

VDA Management

This section describes how to manage your VDA (**syac**) controllers. VDA controllers handle the asynchronous terminal lines on your system.

There are no **sysadm** operations for managing VDA controllers. Instead, you use the **tcload**(1M) command.

To start specific asynchronous terminal controllers, specify the device node from the **/dev/async** directory. For example:

```
# tcload  "/dev/async/syac@60(60000000)"
```

To start all asynchronous terminal controllers, use this command line:

```
# tcload  -a
```

To reset a specific VDC cluster box without resetting all lines on the associated VDA controller, turn the power to the cluster box off and then on. When power returns to the cluster box, the system downloads the required code to the cluster box without user intervention.

End of Chapter

Chapter 12
Printer Management

This chapter tells how to set up and manage printers and print queues using the LP print service. First, the chapter covers the operations that you perform using the **sysadm** utility's Printer menu, which appears in the Device menu.

Second, the chapter covers expert material, the shell-level commands and files that you may use to manage the LP service. The **sysadm** Printer menu operations are implemented using the commands and files discussed in the expert section. The expert section includes troubleshooting tips.

For information on using the **lp** command, see the manual page for **lp**(1).

Printing Now

If you simply want to get a local line printer or local PostScript laser printer (serial or parallel) up and running now, follow these steps:

1. Make sure the printer is properly connected to your computer and powered on. See your hardware documentation.
2. Make sure your DG/UX kernel knows that you have a printer device on your system. If your printer uses a standard interface, such as **lp()** or **syac()**, that is included in your kernel, you may assume that the printer is accessible. If your printer is not a standard device or is connected to a controller that your kernel may not recognize, you need to rebuild the kernel. See Chapter 4 for information on building a kernel.
3. Find out which device pathname (such as **/dev/tty04**) is associated with your printer. You may have to refer to your hardware documentation or your installation manual for a discussion of ports, terminal lines, and device names.
4. Invoke **sysadm** and select the operation `Device -> Printer -> Devices -> Add`. Answer the queries as follows:

   `Printer name`
   Enter a name for the printer. The name may be any arbitrary name that is easy to type and remember.

   `Connection type`
   Select `Local /dev/tty or /dev/lp`.

   `Quick add using default values`
   Select `Yes`.

5. This invokes the quick Add Local Printer operation. Answer the queries as follows:

`Printer type`
: Select the type of printer you are adding from the list provided. This list is generated from the **terminfo** source files in **/usr/src/cmd/terminfo**.

  For a generic 80-column line printer, choose `printer-80`. For a PostScript printer connected to a parallel port, termserver, or a MiLAN Fastport, select `PS-b` for batch mode. For a PostScript printer connected directly to a serial port, choose `PS`.

`Input Types`
: Select the input type or types from the list provided. This list is generated from **/etc/lp/sysadm/input_types**.

  For a generic line printer, choose `simple`. For a PostScript printer, choose only `PS`.

`Device`
: Enter the pathname that you found in Step 3, above.

`Printer description`
: Enter any information you wish to record about the printer. Your comment may be no longer than 40 characters.

Execute the operation. The printer is added to your system.

6. Execute the operation `Device -> Printer -> Devices -> Default`, and make the new printer the default printer.

Now you should be able to print using the **lp**(1) command:

```
# lp myfile ↵
```

You can use similar procedures to quickly add remote printers to your system. For a complete discussion of the queries in the Add operation, see "Adding Devices" in the next section.

LP Management Procedures

This section discusses the operations that **sysadm** offers for managing the LP service. For information on using **sysadm**, invoke **sysadm** and see the Help menu.

To set your LP service up for the first time, start by adding devices. See "Managing Devices." In the **sysadm** Printer menu, select the Devices menu. Before adding any remote printers, you need to add entries for remote hosts using the Systems menu.

The Printer menu comprises the following selections:

`Devices` This menu contains operations to add, delete, modify, and list entries for the printers on your system. These entries determine not only what your printers are called and how the LP service should communicate with them, they also determine other useful features such as what kinds of jobs the printers can accept, how the LP service should handle a printer fault, and so on. You can also name a printer on your system to be the

default printer. You can set the LP service to accept or reject requests submitted to a given printer or class, and you can enable and disable individual printers.

`Classes` This menu contains operations to create, modify, remove, and list status of printer classes.

`Filters` This menu contains operations to add, modify, delete, restore, and list filter programs for use with printers on your system.

`Forms` This menu contains operations to create, modify, remove, and list definitions for pre-printed forms used on your printers. You can also mount forms on printers, which means configure a printer to print that form, and unmount forms from printers. In addition, you can display the current mount status of defined forms.

`Priorities`
This menu contains operations to set the default priority for print requests. You can also set and remove priorities for individual users, and you can set a default priority for print requests. A list operation displays current priority settings.

`Requests` This menu contains operations to list and cancel job requests as well as to hold requests and to resume held requests. There is also an operation for moving requests from one printer to another.

`Scheduler/Service`
This menu contains operations to start and stop the LP service scheduler and to display the current status of the scheduler.

`Systems` For systems on a network, this menu contains operations to add, modify, remove, and list entries for remote systems with which you can share printer services.

`List` This operation displays the status of the LP service, including LP service database information about the default printer setting, devices, classes, remote systems, forms, filters, and current status of the scheduler, devices, classes, forms, and jobs requests.

You need to be **root** or **lp** (uid 0 or 6) to use these operations, except for the Requests menu operations that cancel, hold, and resume print jobs. Except for the List operation in the Filters menu, you do not need to be superuser to invoke any of the List operations.

The following sections discuss the Printer menu operations and LP services concepts in more detail.

Managing Devices

Before you can perform any printer-related operations, you need to define the printers that are available on the system. If you are adding printers connected to the local system, use the operations in the `Device -> Printer -> Devices` menu. If you are adding printers connected to other systems in your network, first add the remote systems with the `Device -> Printer -> Systems -> Add` operation before proceeding to the `Device -> Printer -> Devices` menu.

This section discusses the various operations that **sysadm** offers for printer device management:

`Add`	This operation adds a device entry to the LP service databases.
`Default`	This operation allows you to define a printer for requests that do not specify a destination.
`Delete`	Use this operation to remove the file, directories, and database entries that support a particular printer.
`Modify`	Use this operation to change the attributes of an existing printer.
`Accept`	This operation allows a printer to accept print requests.
`Reject`	This operation makes the LP service refuse requests for the printer. If a user submits a job for printing at this printer, the **lp** command returns an error.
`Disable`	This operation causes the LP service to cease sending jobs to a printer.
`Enable`	This operation tells the LP service to verify that a printing device is functional and if so to make it available for printing.
`List`	This operation produces a general status listing for your system's printing devices.

The following sections discuss these operations in more detail.

Adding Devices

If you are adding a remote printer or printer class, you need to verify that you have already added an entry for the remote system using the following operation `Device -> Printer -> Systems -> Add`.

After you have added the system entry, use the Add operation documented here to set up access to the printer.

To declare a printer device to the LP service, use the Add operation in the Devices menu. This operation not only sets up the files and directories that the LP service needs to use the printer, it also lets you customize the printer's behavior.

You can add a printer to your system in two ways: Quick Add or Full Add. If you want to get a printer up and running immediately, you should choose Quick Add. Quick Add minimizes the number of prompts you have to answer by using established defaults for most of a printer's characteristics. If you choose Full Add, you must specify all of a printer's characteristics.

When you select `Device -> Printer -> Systems -> Add`, you must answer the following three queries:

- `Printer name`
- `Connection type`
- `Quick Add using default values`

The following sections discuss these queries in more detail.

`Printer name`
A printer name may be any name that you choose. You should choose a name that tells something about the printer, such as what type of printer it is, its location at your workplace, or its location on your network. It is also a good idea to choose a name that is short and easy to remember.

`Connection type`
This is the type of printer that you want to add. There are five different connection types:

`Local /dev/tty or /dev/lp`
A printer connected to your system through a device whose name begins with `/dev`.

`Remote lpNet to DG/UX 5.4 system`
A remote printer on a DG/UX Release 5.4 or later system running the **lpNet** daemon.

`Remote lpNet to lpd system`
A remote printer on a system that supports the Line Printer Daemon Protocol (**lpd**), also known as the Berkeley print spooler.

`Remote to system other than DG/UX 5.4`
A remote printer on a system that is not running DG/UX Release 5.4 or later. The following systems are supported:

remshlp: a remote UNIX host which supports the **lp** command using **remsh**.

remshlp_bsd: a remote UNIX host which supports the **lpr** command using **remsh**.

remshlp_a: a remote AOS/VS II host with TCP/IP II installed which supports the **lp** command using **remsh**.

`Remote networked printer device`
A remote printer connected to a network device independent of a host. The following devices are supported:

termprinter: used for termserver connections.

fastport: used for MiLAN Fastport networked printer servers with a serial and parallel port.

`Quick Add using default values`
Select this option if you want to use Quick Add to add a printer to your system. Quick Add is covered in more detail in the next section. If you do not select this option, you have to do a Full Add, which is covered in a later section.

Quick Add

The Quick Add operation enables you to get a printer up and running immediately. Quick Add minimizes the number of queries you have to answer for each connection

type by taking some standard defaults for the majority of a printer's characteristics. Each connection type, except for the two **lpNet** connection types, has its own set of queries for adding a printer. See "Printing Now" at the beginning of this chapter for an example of adding a local printer to your system. *Customizing the DG/UX™ System* has additional examples of Quick Add.

The following sections cover the Quick Add queries and the defaults associated with each connection type. For a complete description of each printer characteristic, see "Full Add."

Local /dev/tty or /dev/lp

You must answer the following queries when you use Quick Add to add a local printer to your system:

`Printer type`

Select a printer type. A printer's type refers to a **/usr/src/cmd/terminfo** database entry that describes control codes the printer requires for handling various initialization and configuration operations as well as printing attributes. You choose from a list generated from the **terminfo** database.

If you customize or add a new **terminfo** entry for a printer, either include the source code in the **/usr/src/cmd/terminfo** directory or use the **tic**(!M) command to create a compiled version for the **/usr/lib/terminfo/...** directory.

`Input types`

Select one or more input types, also called content types, for every printer you add. An input type is a designation telling what kind of files you can print on the printer. You choose from a list generated from the **/etc/lp/sysadm/input_types** file.

`Device`

Enter a device pathname. The pathname, indicating an entry in the **/dev** directory, must already exist.

`Description`

This is a comment field for entering any information you wish to record about the printer. The field may contain no more than 40 characters.

Quick Add for a local printer assumes the following defaults:

`Interface script`	standard
`Stty options`	none
`Input options`	none
`Device is also a login terminal`	no
`Fault recovery`	continue
`Alert messages`	mail
`Alert interval`	0
`Users to allow`	all

User to deny	N/A
Forms to allow	none
Forms to deny	N/A
Enable printer	yes
Accept printer	yes

Remote lpNet to DG/UX 5.4 system and Remote lpNet to lpd system
You must answer the following queries when you use Quick Add to add an **lpNet** connected remote printer to your system:

Remote system
Enter the name of the remote system to which the printer is attached. See the **lpsystem**(1M) manual page and the "Expert Information" section later in this chapter for more information on the name.

Remote printer name
Enter the name of the printer on the remote system.

Description
This is a comment field for entering any information you wish to record about the remote printer. The field may contain no more than 40 characters.

Quick Add for a remote **lpNet** printer assumes the following defaults:

Printer type	printer-80
Input types	simple
Users to allow	all
User to deny	N/A
Forms to allow	none
Forms to deny	N/A
Enable printer	yes
Accept printer	yes
Interface script	N/A
Stty options	N/A
Input options	N/A
Device is also a login terminal	N/A
Fault recovery	N/A
Alert messages	N/A
Alert interval	N/A

Remote to system other than DG/UX 5.4
You must answer the following queries when you use Quick Add to add a remote printer to your system from a non-DG/UX system or pre-5.4 system using **remshlp**:

`Printer type`
Select a printer type. A printer's type refers to a **/usr/src/cmd/terminfo** database entry that describes what control codes the printer requires for handling various initialization and configuration operations as well as printing attributes.

For a generic 80-column line printer, choose `printer-80`. For a PostScript printer connected to a parallel port, termserver, or a MiLAN Fastport, select `PS-b` for batch mode. For a PostScript printer connected directly to a serial port, choose `PS`.

`Input types`
Select one or more input types, also called content types, for every printer you add. An input type is a designation telling what kind of files you can print on the printer.

For a generic line printer, choose `simple`. For a PostScript printer, choose only `PS`.

`Interface scripts`
Select the remote interface script, which is the interface script the LP service uses to drive the printer. You choose from a list generated from the **/etc/lp/sysadm/remote_models** file. The provided scripts all use the **Remsh** model.

`Remote system`
Enter the name of the remote system to which the printer is attached. See the **lpsystem**(1M) manual page and the "Expert Information" section later in this chapter for more information on the name.

`Remote printer name`
Enter the name of the printer on the remote system.

`Description`
This is a comment field for entering any information you wish to record about the remote printer. The field may contain no more than 40 characters.

Quick Add for a remote printer assumes the following defaults:

`Stty options`	none
`Input options`	none
`Fault recovery`	continue
`Alert messages`	mail
`Alert interval`	0
`Forms to allow`	none
`Forms to deny`	N/A
`Users to allow`	all
`User to deny`	N/A
`Enable printer`	yes
`Accept printer`	yes
`Device is also a login terminal`	N/A
`Device`	N/A

Remote networked printer device
You must answer the following queries when you use Quick Add to add a networked printer, such as a termserver or MiLAN Fastport, to your system:

`Printer type`
Select a printer type. A printer's type refers to a **/usr/src/cmd/terminfo** database entry that describes what control codes the printer requires for handling various initialization and configuration operations as well as printing attributes.

For a generic 80-column line printer, choose `printer-80`. For a PostScript printer connected to a parallel port, termserver, or a MiLAN Fastport, select `PS-b` for batch mode. For a PostScript printer connected directly to a serial port, choose `PS`.

`Input types`
Select one or more input types, also called content types, for every printer you add. An input type is a designation telling what kind of files you can print on the printer.

For a generic line printer, choose `simple`. For a PostScript printer, choose only `PS`.

`Interface scripts`
Select the network interface script, which is the interface script the LP service uses to drive the printer. You choose from a list generated from the **/etc/lp/sysadm/remote_models** file.

`Networked printer device name`
Enter the name of the networked printer device. You can get the name from **/etc/hosts** or through the **ping**(1M) command.

`Port number or name`
Optionally, you may enter the port number or name for the networked printer device. The default port number is decimal 23.

`Description`
This is a comment field for entering any information you wish to record about the networked printer. The field may contain no more than 40 characters.

Quick Add for a networked printer assumes the following defaults:

`Stty options`	none
`Input options`	none
`Fault recovery`	continue
`Alert messages`	mail
`Alert interval`	0
`Users to allow`	all
`User to deny`	N/A
`Forms to allow`	none
`Forms to deny`	N/A

```
Enable printer          yes
Accept printer          yes
Device                  N/A
Device is also a login terminal    N/A
```

Full Add

The Full Add operation gives you more control over the characteristics of the printer you are adding than the Quick Add operation. However, Full Add requires you to specify each of a printer's characteristics which can be somewhat lengthy. If you want to get a printer up and running quickly, you should use the Quick Add operation.

As with the Quick Add operation, the queries for adding a printer vary with the connection type of the new printer. The following sections discuss the queries for each connection type.

Local /dev/tty or /dev/lp

You must answer the following queries when you use Full Add to add a local printer to your system:

`Printer type`

A printer's type refers to a **terminfo** database entry that describes what control codes the printer requires for handling various initialization and configuration operations as well as printing attributes.

Initially, the default printer type is **printer–80**. Although you may leave the printer's type set to the default, you can enhance your printer's ability to serve users by assigning it a more appropriate type.

The following list contains some accepted printer types.

printer	Generic line printer; 132 columns.
printer-80	Generic line printer; 80 columns.
PS	Serial PostScript printer, normal paper stacking.
PS-r	Serial PostScript printer, reverse order paper stacking.
PS-b	Parallel or TermServer PostScript printer, normal paper stacking.
PS-br	Parallel or TermServer PostScript printer, reverse order paper stacking.
citoh	Citoh 8510 printer.
daisy	Daisy brand printer.
qume	Qume Sprint 11.
la100	DEC LA 100 printer.
ln03	DEC LN03 laser printer.
epsonfx	Epson FX printer; 136 columns.
epsonfx-80	Epson FX printer; 80 columns.

If your printer can emulate multiple printers, you may specify more than one type. The **terminfo** entries for emulated printer types generally have names of the form

printer–emulation, where *printer* is the model of printer and *emulation* is the type of printer that it emulates. For example, **dg6617–epsonfx** is a Data General model 6617 printer emulating an Epson printer.

`Input types allowed`
You may set one or more input types, also called content types, for every printer you add. An input type is a designation telling what kind of files you can print on the printer. Initially, the default input type is **simple**, which means common ASCII character text. PostScript printers are restricted to type `PS` only.

Input types also affect which filters are used. See the "Managing Filters" section later in this chapter for more information.

Generally, most printers can print **simple** files and files whose input type is the same as the printer type. Additionally, you can make up other input types to correspond to printer types and the types of files you print. Here are some accepted input types:

`cif`	Output of BSD `cifpbt` processor.
`epsonfx`	Epson and compatible printers.
`proprinter`	IBM ProPrinter™ and compatible printers.
`fortran`	ASA carriage control format.
`laserjet`	Hewlett-Packard LaserJet® and compatible printers.
`otroff`	CAT typesetter instructions generated by BSD (old) `troff`.
`plot`	Plotting instructions for Tektronix displays and devices.
`PS`	PostScript language.
`raster`	Raster bitmap format for Varian raster devices.
`simple`	ASCII file.
`tex`	DVI format files.
`troff`	Device-independent output from the `troff` text-formatting processor.

Once you have assigned input types to your printers, users may submit requests specifying the input type with the **lp** command's **–T** option. As long as the request does not specify a destination and there is no default printer, the LP service will print the request on any printer that accepts the specified type. If a default printer has been set and you want to override it, specify **–d any** on the **lp** command line.

For example, the following command line overrides the system default printer and submits the PostScript file **myfile** to any printer that accepts input type `PS`:

```
% lp -d any -T PS myfile ↵
```

Note that **myfile** must already be in PostScript format. In general, when using the **lp** option **–T** *type*, the file should be in the specified *type* already.

`Interface script`

Enter the interface script that the LP service uses to drive the printer. For a common line printer or laser printer not falling into one of the categories below, select the **standard** model. Other models are:

dg455x	For Data General text-only laser printers, models 4557 and 4558, or the replacement model 6454.
fastport	For networked printers using the Milan network printer box.
remshlp	For remote printers using a System V LP scheduler (the DG/UX default scheduler).
remshlp_bsd	For remote printers using a BSD spooler.
remshlp_a	For remote printers on an AOS/VS system.
termprinter	For TermServer printers.
standard	For local line and PostScript printers.

You may copy these interface scripts, located in **/var/spool/lp/model** (links to **/usr/lib/lp/model**), and tailor them to fit printers on your system. For more information about interface scripts, see "Expert Information" later in this chapter.

`Device`

You must supply a device pathname when adding a local printer. The pathname, indicating an entry in the **/dev** directory, must already exist.

The device pathname represents the port to which you have connected the printer. If your computer hardware has a dedicated line printer port, this pathname is **/dev/lp**.

NOTE: If you have connected your printer to an asynchronous terminal line, you should refer to your installation documentation and/or your controller documentation to ascertain which line it is. An asynchronous terminal pathname has the form **/dev/tty***nn*, where *nn* is the line number.

If you have connected your printer to an asynchronous line, do not assign a port monitor service to that line. Use the Ports menu operations, making sure any service for the line is disabled.

`Stty options`

Specify additional **stty** options in this query. **Stty** options control line I/O characteristics such as baud, flow control, parity, output processing, and so on. For binary printing, use the **stty** option **–opost**. For more information, see the **stty**(1) manual page.

`Print options`

You may specify several print options to determine how the printer spaces jobs on the page. These options are:

`length` This option determines the number of lines printed per page. For example, to set page length to 63 lines, include the following option:

```
length=63
```

`width` This option determines the width of the printable area on the page. Specify the width in characters. For example, to set page width to 40 characters, include the following option:

```
width=40
```

`cpi` This option determines the horizontal printing density, specified in characters-per-inch. For example, to set characters-per-inch to 10, specify the following option:

```
cpi=10
```

`lpi` This option determines the vertical printing density, specified in lines-per-inch. For example, to print 6 lines per inch, include the following option:

```
lpi=6
```

`banner/nobanner/banneroff`
The banner is the job header page. By default, the LP service does not print a banner before every job. **banneroff** means never print a banner page. To allow users to request that a job be printed with a banner, specify the following option:

```
banner
```

When specifying these attributes for the Print Options query, list them separated by spaces, for example:

```
length=63 lpi=6
```

To set any of these printing attributes to their default values, specify them without a number, for example:

```
cpi=
```

In the case of the `banner` attribute, reset the default by specifying `banner`.

`Device is also a login terminal`

In some cases, the device used as a printing output device may also be used for logging into the system.

`Fault recovery`

The LP service can handle a printer fault in one of three ways:

`continue` When the printer encounters a fault, it either waits for the fault to clear and reprints the current page or waits for the fault to

clear and resumes printing at the point the fault occurred. The particular action taken depends on the capability of the printer, and requires a special filter not included with the supplied **lp** subsystem. Without the special filter, the print job is reprinted from the beginning after the fault clears and the printer is enabled.

beginning When the printer encounters a fault, it clears the fault condition and reprints the entire job.

wait When the printer encounters a fault, the LP service stops the job and places the printer in a disabled state. The printer will not continue until you enable it. Enable a printer with the **sysadm** operation Device –> Printer –> Enable.

`Alert messages`

When a printer fault occurs, the LP service can alert you to the condition in several different ways:

mail Using the **mail**(1) command, the LP service sends mail to the administrator.

write Using the **write**(1) command, the LP service writes a message to the terminal where the administrator is currently logged in.

quiet Send no alert for the current print fault. You can select this option only if the LP service has already determined that your printer is in a fault state. This option is useful for silencing alerts for a fault condition that you already know exists.

none The LP service sends no alerts if it detects a fault.

As an alternative, you may write your own alert script. The script should accept the alert message as standard input from the LP service. When prompted, specify the command line for invoking your custom alert script.

You may also specify how often the LP service alerts you when a fault occurs. You specify the interval in minutes. For example, an interval of 10 means that the LP service will send out an alert every 10 minutes for as long as the fault condition exists. If you want the LP service to send out only one alert message per fault, specify 0.

`Alert interval`

When a printer fault occurs, the system sends periodic messages to alert the system administrator. The alert interval determines the number of minutes that the LP system will wait between messages. Specifying **0** indicates that you want only one message per fault sent to the administrator.

`Users to allow`

Optionally, you may specify which users may use the printer. Regardless of any users expressly denied access to the printer (via `Users to deny`, discussed below), only the users specified in `Users to allow` may use the printer. Specifying **all** allows access for all users except any who are denied access in `Users to deny`.

`Users to deny`
Optionally, you may specify which users may not use the printer. If you specify users in `Users to allow`, discussed previously, it is unnecessary to specify users in `Users to deny`.

`Forms to allow`
Specify forms, added with the Create operation of the Forms menu, that can be used on this printer. Regardless of any forms expressly restricted from the printer (via `Forms to deny`, discussed below), only the forms specified in `Forms to allow` may be used on the printer. Specifying **all** allows use of all forms except any that are restricted in `Forms to deny`.

`Forms to deny`
Specify forms that may not be used on this printer. If you specify forms in `Forms to allow`, discussed previously, it is unnecessary to specify forms in `Forms to deny`.

`Enable`
Enabling the printer allows the LP service to send waiting jobs to it. If a printer is set to accept but it is disabled, the LP service accepts jobs submitted for the printer, but the LP service will not send the jobs to the printer for printing.

The Device menu provides operations for enabling and disabling printers.

`Accept`
Setting the printer to accept requests allows users to submit jobs for that printer. If a user submits a job for a printer that is set to reject jobs, the **lp** command does not queue the job; instead, it returns a reject message to the user.

The Device menu provides operations for setting printers to accept and reject job requests.

`Printer Description`
This is a comment field for recording any information you wish to record about the printer. The field may contain no more than 40 characters.

Remote lpNet to DG/UX 5.4 system and Remote lpNet to lpd system
When adding a remote printer, make sure that the printer is physically attached to the remote system you specify. The DG/UX system does not allow you to send print requests through one remote system to reach a printer attached to yet another remote system.

You must answer the following queries when you use Full Add to add an **lpNet** connected remote printer to your system. For complete information about queries already covered, see the descriptions above for local printers.

`Printer type`
A printer's type refers to a **terminfo** database entry that describes what control codes the printer requires for handling various initialization and configuration operations as well as printing attributes.

`Remote system`
Enter the name of the remote system to which the printer is attached.

`Remote printer name`
Enter the name of the printer on the remote system.

`Users to allow`
Optionally, you may specify which users may use the printer.

`Users to deny`
Optionally, you may specify which users may not use the printer. If you specify users in `Users to allow`, discussed above, it is unnecessary to specify users in `Users to deny`.

`Forms to allow`
Specify forms, added with the Create operation of the Forms menu, that can be used on this printer.

`Forms to deny`
Specify forms that may not be used on this printer. If you specify forms in `Forms to allow`, discussed above, it is unnecessary to specify forms in `Forms to deny`.

`Enable`
Enabling the printer allows the LP service to send waiting jobs to it.

`Accept`
Setting the printer to accept requests allows users to submit jobs for that printer.

`Printer Description`
This is a comment field for recording any information you wish to record about the printer.

Remote to system other than DG/UX 5.4
When adding a remote printer, make sure that the printer is physically attached to the remote system you specify. The DG/UX system does not allow you to send print requests through one remote system to reach a printer attached to yet another remote system.

You must answer the following queries when you use Full Add to add a remote printer from a non-DG/UX system or pre-5.4 system using **remshlp** to your system. For complete information about queries already covered, see the descriptions above for local printers.

`Printer type`
A printer's type refers to a **terminfo** database entry that describes what control codes the printer requires for handling various initialization and configuration operations as well as printing attributes.

`Input types allowed`
You may set one or more input types, also called content types, for every printer you add. PostScript printers are restricted to `PS` only.

`Interface script`
Enter the interface script that the LP service uses to drive the printer.

`Remote system`
Enter the name of the remote system to which the printer is attached. See the **lpsystem**(1M) manual page and the "Expert Information" section later in this chapter for more information on the name.

`Remote printer name`
Enter the name of the printer on the remote system.

`Stty options`
Specify additional **stty** options in this query. For binary printing, use the **stty** option **–opost**. For more information, see the **stty**(1) manual page.

`Print options`
You may specify several print options to determine how the printer spaces jobs on the page.

`Fault recovery`
Specify how the LP service handles a printer fault.

`Alert messages`
Specify how the LP service alerts you if there is a printer fault.

`Alert interval`
Specify the amount of time between alert messages sent when a printer fault occurs.

`Users to allow`
Optionally, you may specify which users may use the printer.

`Users to deny`
Optionally, you may specify which users may not use the printer. If you specify users in `Users to allow`, discussed above, it is not necessary to specify users in `Users to deny`.

`Forms to allow`
Specify forms, added with the Create operation of the Forms menu, that can be used on this printer.

`Forms to deny`
Specify forms that may not be used on this printer. If you specify forms in `Forms to allow`, discussed above, it is not necessary to specify forms in `Forms to deny`.

`Enable`
Enabling the printer allows the LP service to send waiting jobs to it.

`Accept`
Setting the printer to accept requests allows users to submit jobs for that printer.

`Printer Description`
This is a comment field for recording any information you wish to record about the printer.

Remote networked printer device
When adding a remote printer, make sure that the printer is physically attached to the remote system you specify. The DG/UX system does not allow you to send print requests through one remote system to reach a printer attached to yet another remote system.

You must answer the following queries when you use Full Add to add a networked printer, such as a termserver or MiLAN Fastport, to your system. For complete

information about queries already covered, see the descriptions above for local printers.

`Printer type`
A printer's type refers to a **terminfo** database entry that describes what control codes the printer requires for handling various initialization and configuration operations as well as printing attributes.

`Input types allowed`
You may set one or more input types, also called content types, for every printer you add. PostScript printers are restricted to **PS** only.

`Interface script`
Enter the interface script that the LP service uses to drive the printer.

`Networked printer device name`
Enter the name of the networked printer device, such as from **/etc/hosts**.

`Port number or name`
Optionally, you may enter the port number or name for the networked printer device.

`Stty options`
Specify additional **stty** options in this query.

`Print options`
You may specify several print options to determine how the printer spaces jobs on the page.

`Fault recovery`
Specify how the LP service handles a printer fault.

`Alert messages`
Specify how the LP service alerts you if there is a printer fault.

`Alert interval`
Specify the amount of time between alert messages sent when a printer fault occurs.

`Users to allow`
Optionally, you may specify which users may use the printer.

`Users to deny`
Optionally, you may specify which users may not use the printer. If you specify users in `Users to allow`, discussed above, it is not necessary to specify users in `Users to deny`.

`Forms to allow`
Specify forms, added with the Create operation of the Forms menu, that can be used on this printer.

`Forms to deny`
Specify forms that may not be used on this printer. If you specify forms in `Forms to allow`, discussed above, it is not necessary to specify forms in `Forms to deny`.

`Enable`
Enabling the printer allows the LP service to send waiting jobs to it.

`Accept`
Setting the printer to accept requests allows users to submit jobs to that printer.

`Printer Description`
This is a comment field for recording any information you wish to record about the printer.

Adding Printers with Multiple Emulation Modes

The DG/UX system provides printer specific terminfo entries that allows one physical printer to be used in three different emulation modes. These terminfo entries contain two emulation modes per printer type: PostScript (PS) and Printer Control Language (PCL). In addition, this file contains entries for PostScript printers that are connected over a parallel port (PS–b). The printer specific terminfo entries are listed as follows;

hplaserjet4–PS – Laserjet 4 PostScript (local serial connection)
hplaserjet4–PS–b – Laserjet 4 PostScript (any non local serial)
hplaserjet4–PCL – Laserjet 4 PCL mode

hplaserjet3Si–PS – Laserjet 3Si PostScript (local serial connection)
hplaserjet3Si–PS–b – Laserjet 3Si PostScript (any non local serial)
hplaserjet3Si–PCL – Laserjet 3Si PCL mode

hplaserjet3–PS – Laserjet 3 PostScript (local serial connection)
hplaserjet3–PS–b – Laserjet 3 PostScript (any non local serial)
hplaserjet3–PCL – Laserjet 3 PCL mode

dg6640T–PS – DG Model 6640T printer in PostScript (local serial)
dg6640T–PS–b – DG Model 6640T printer in PostScript (any non serial)
dg6640T–hp2 – DG Model 6640T printer in PCL mode

dg6646T–PS – DG Model 6646T printer in PostScript (local serial)
dg6646T–PS–b – DG Model 6646T printer in PostScript (any non serial)
dg6646T–hp2 – DG Model 6646T printer in PCL mode

dg6771–PS – DG Model 6771 printer in PostScript (local serial)
dg6771–PS–3 – DG Model 6771 printer in PostScript (any non serial)
dg6771–hp3 – DG Model 6771 printer in PCL mode

dg6772–PS – DG Model 6772 printer in PostScript (local serial)
dg6772–PS–b – DG Model 6772 printer in PostScript (any non serial)
dg6772–hp2 – DG Model 6772 printer in PCL mode

Figure 12–1 Printer–Specific Terminfo Entries

Use **sysadm** to add a printer with multiple emulation modes. For example, to add a single printer to operate in hplaserjet4–PS and hplaserjet4–PCL emulation modes, peform the operation `Device -> Printer -> Devices -> Add` twice. Each time, you must specify a different printer name, a different (desired) emulation type, but the same printer device.

For example:

Sysadm Operation	Printer Name	Printer Type	Device
Add first printer	foo	hplaserjet4–PS	/dev/tty/5
Add second printer	bar	hplaserjet4–PCL	/dev/tty/5

When you submit a print request to a printer, its name is associated with the appropriate emulation sequences. For example, the print command, **lp –dfoo tax–file**, invokes the hplaserjet4–PS emulation mode initialization sequence.

Setting the Default Printer

You can make one of your printers the default printer so that any job submitted without the destination (**–d** *destination*) option goes automatically to the default printer. The default printer may be either a local printer or a remote one.

Deleting Devices

The Delete operation removes the files, directories, and LP service database entries supporting the specified printer. Deleting an entry for a remote printer has no effect on the files, directories, and LP service database entries supporting the printer on the remote system. Deleting the last printer of a class also removes the class.

Modifying Devices

Use the Modify operation to change the attributes that you set for a printer when you added it. For a complete discussion of the various attributes, see "Adding Devices." The Modify operation for printers uses the same queries as a Full Add.

The Modify operation does not allow you to change the class to which a device belongs. To change the class membership, use the Modify operation of the Classes menu.

Accepting or Rejecting Requests

Use the Accept operation to allow users to submit requests to a printer device. To make the LP service reject requests for a given printer, use the Reject operation.

When a user submits a request to a printer that is rejecting requests, the LP service returns an error.

Setting a printer to reject requests does not remove jobs for that printer that are already in the print queue.

The accept or reject status of a printer has no effect on its enable or disable status, described below.

Enabling and Disabling Devices

The Enable operation makes a printer available to print requests. If the fault recovery attribute for a printer is set to "wait" (see "Adding Devices"), you will have to enable a printer after a fault occurs.

The Disable operation makes a printer unavailable to print jobs. Users may still submit jobs for printing by the printer even though it is disabled (assuming it is set to accept requests). Requests submitted for a disabled printer will remain in the print queue until you remove them, move them, or enable the printer.

The Disable operation allows you to select how to handle a job currently being printed by the printer. You may choose to interrupt the job, in which case the operation cancels the job, removing it from the queue, or you may choose to let the printer finish the request before entering the disabled state.

When you disable a printer, you may specify a reason in the form of a line of text of your own choosing, that tells why you have disabled the printer. When a user displays the status of the printer, the display includes the reason that you entered.

Displaying Devices

Use the List operation to display information about any or all printers that you have added on your system. When you invoke the List operation, you may select either of two kinds of listing:

`setup` This listing shows not only some information about the current state of the printer, but it also displays the settings of attributes that you defined for the printer when you created it with the Add operation:

Forms mounted
Content types
Printer types
Description
Connection
Interface
On fault
After fault
Users allowed
Forms allowed
Banner
Character sets
Default pitch
Default page size
Default port settings

See "Adding Devices" for a discussion of these attributes.

`state` This listing shows only information pertinent to the current state of the printer, whether it is enabled and accepting requests, and how many requests are in its queue. The display also includes the device pathname.

Managing Classes

A printer class is a group of printers to which you may submit a print request for printing by any printer in the class. When you specify a class on the **lp** command line, the LP service prints the job on the first available printer in the class. The advantage to this system is that users do not need to investigate the print queue every time they submit a job in order to see which printer is currently available or which printer already has jobs queued up. The result is faster throughput for your users and more efficient utilization of your printers.

Another advantage of print classes is that they allow you to maximize use of favored printers. For example, if you have a fast line printer and a slow line printer, you can create a class where the fast line printer is first and the slow one is second. Thus, the LP service submits the job to the fast printer if both are available and to the slow printer only when the fast printer is already in use.

Creating Classes

The Create operation of the Classes menu creates a new printer class and allows you to add printers to the class. When you create the class, you must add at least one printer to it. The printers you add should already exist; add printers using the operation `Device -> Printer -> Devices -> Add`.

A class may contain only local printers. See "Managing Devices" in this chapter for more information.

When you create a class, the order in which you list the printers determines the order in which the LP service checks them when assigning jobs. When you submit a job to the class, the LP service assigns the job to the first printer in the list that is available. A printer may belong to more then one class.

The Create operation also allows you to specify whether you want the class to accept jobs or reject jobs. For normal operation, specify Accept. If you do not want users to use the class at this time, specify Reject. Even if you set a class to reject requests, a user can still use a printer in the class by specifying the printer name on the **lp** command line (using the **–d** option).

To submit a job for printing by a certain class, use the **–d** option of the **lp** command. For example, to print the file **myfile** at the first available printer in class **class1**, use this command:

```
% lp -d class1 myfile ↵
```

Modifying Classes

Use the Modify operation of the Classes menu to modify a printer class. The class must already exist.

The Modify operation allows you to add printers to the class or remove members from the class. If you add a member to the class, the printer must already exist. Add a printer with the operation `Device -> Printer -> Devices -> Add`.

To change the order of members in the class, you need to remove them and add them again in the desired order. The order of printers in the class determines the order in which the LP service checks them when assigning jobs. When you submit a job to the class, the LP service assigns the job to the first printer in the list that is available.

While you are adding or removing printers, you can also change the accept or reject status of the class. Even if you set a class to reject requests, a user can still use a printer in the class by specifying the printer name on the **lp** command line (using the **–d** option).

Removing Classes

Use the Remove operation to delete a printer class from the LP service. You do not have to delete the printers from the class before removing the class.

Displaying Classes

To display information about the classes on your system, select the List operation from the Classes menu. You may list the status of all classes on the system, or you may specify a class for listing. The operation lists the printers that are members of the classes.

Managing Filters

Filters are programs that process files before printing them. This section covers the **sysadm** operations, located in the Filters menu of the Printer menu, that you use to manage LP service filters.

A filter can function in three ways:

- To convert a file from one format to another before printing.
- To handle the kind of special printing modes that a user can request with the **lp** command's **–y** option: two-sided printing, landscape printing, draft quality printing, and so on.
- To detect printer faults and notify the LP service.

When defining a filter, you not only name the program that functions as the filter, you also specify what content types the filter will accept as input and produce as output.

A content type refers to the formatting, codes, or conventions used in a file to describe its page layout or contents. The purpose of having content types is to allow the LP service a means of matching user-submitted print requests with compatible filters and printers. When a user submits a job and specifies a content type, the LP service attempts to match the file's content type with the printer type or input type specified for a printer. If no satisfactory printer exists, the LP service looks for filters that can convert the file into a printable content type.

For example, you may have created filters that accept or produce PostScript content. If you submit a file for printing and you specify that the file is of type PostScript, the LP service will look for a printer whose input type is PostScript. If the print service does not find one, it will attempt to assemble a series of filters that can take your PostScript file and convert it to a type that one of your printers can print.

You set a printer's input content type when you add it with the Add operation in the Devices menu. To change a printer's input types, use the operation `Device -> Printer -> Devices -> Modify`.

The Filters menu provides operations for:

- Adding filters.
- Modifying filters.
- Removing filters.
- Restoring filters.
- Displaying filters.

The following sections elaborate on these operations in more detail.

For information on writing a filter, see the section "Providing Filters" in the "Expert Information" section of this chapter.

Adding Filters

Use the Add operation to introduce a filter program to the LP service. When invoking the Add operation, be prepared to supply the following information:

- The command line for invoking the filter.
- The existing filter, if any, to copy to make the new filter.
- The content types the filter will accept as input.
- The content types the filter will produce as output.
- The printer types compatible with the output types.
- The printers that may print the filter's output.
- The filter speed.
- The options to use when invoking the filter.

The following sections elaborate.

Command Line

This is the full pathname of the filter program. If there are any fixed options that the program always needs, include them here.

Filter to Copy

This is any existing filter that you may want to copy to make the new filter.

Input Types

This is the list of file content types that the filter can process. The content type designations are names that you make up based on the types of files your users print and the types of printers you have. The LP service matches file content types and filter output types with filter input types, so be careful to use consistent naming conventions. Most filters can accept only one content type as input.

Output Types

This is the list of file content types that the filter can produce. The LP services matches content output types with filter input types and printer input types, so, again, use consistent naming conventions. Most filters can produce only one content type as output.

Printer Types

This is the list of printer types for which the filter may produce output. In most cases, this list will be the same as the Output Types list. Any listed printer types should match types of existing printer devices.

Printers

There may be printers who are compatible with the filter's output type but for other reasons are undesirable as output devices for this filter. When this case is true, you should specify the desirable printers in the Printers list, omitting the undesirable ones of the same type.

Filter Speed

You can designate a filter as either fast or slow. When a print request requires a filter designated "fast," the print service assigns a printer to the request at the same time it starts the filter. If a filter is particularly slow, however, this implementation occupies the printer unnecessarily because the printer is out of use while it waits for the filter to finish.

To avoid this kind of waste of printer time, you can designate such filters as "slow." The LP service executes slow filters in the background without causing the printer to wait; the slow filter does not access the printer (or the next filter in the series) until it has filtered the entire request. While the slow filter works, other requests are free to print.

If you are adding a filter that is intended to detect printer faults, you must designate the filter as "fast." You may designate a filter as "slow" if it does not require access to the printer.

Slow filters that are invoked by modes (via the **lp** command's **–y** option) must be run on the system where the print request was issued. The LP service cannot pass values for modes to remote systems. It can, however, match a file content type

(specified after the **-T** option) to a content type on a remote system. Therefore, to activate special modes on a remote system, specify content types that will cause the LP service to match input types and output types.

Options

The Options field allows you to specify how user options, LP service database settings, and preceding filter options can determine options to be passed to this filter.

For example, you can set the options field so that if a specified character set is in effect for this print request, the LP service should include a particular option on the filter command line.

Specifically, the Options field allows you to determine the filter invocation command line based on these printing attributes:

- Input content type.
- Output content type.
- Printer type.
- Printer name.
- Character pitch (characters per inch).
- Line pitch (lines per inch).
- Page length.
- Page width.
- Pages to print.
- Character set.
- Form name.
- Number of copies.
- Special modes.

For a detailed discussion on how to set the Options field for these attributes, see "Defining Options with Templates" in "Expert Information."

Modifying Filters

Use the Modify operation to change the definition for an existing filter. For a discussion of the features that make up a filter, see "Adding Filters."

Removing Filters

Use the Remove operation to delete the LP service database entries that support a given filter. The operation does not remove the filter program itself.

Restoring an Original Filter Definition

Use the Restore operation to restore a filter's definition to the definition that originally shipped with the system. This operation functions only for filters that shipped with the DG/UX system.

Displaying Filters

This operation displays the attributes of filters currently available on your system. These attributes include:

`Filter command name`	The command name that invokes the filter.
`Input content type`	The list of content types that the filter accepts as input.
`Output content type`	The list of content types that the filter produces as output.
`Printer type`	The list of printer types for which the filter produces output.
`Printer names`	The list of printers for which the filter produces output.
`Speed`	The filter's speed.
`Command`	Any options that the LP service passes to the filter.

The "Adding Filters" section, earlier in this chapter, discusses these attributes in more detail.

Managing Forms

A form is a description of a page layout, plus some other attributes, that determine how a printer loaded with pre-printed form stock should complete special requests. For example, if you have a printer loaded with pre-printed purchase order forms, the form description provides the printer with descriptive information such as lines per inch, characters per line, required print wheels, ribbon colors, and so on. The form description can also include an alignment pattern that you can print to make sure that the form stock is aligned correctly in the printer.

Once you have defined a form, you associate it with a printer by mounting it on that printer. The LP service then restricts use of the printer to printing requests that require that form. If a user submits a job requiring a form that is not mounted on the desired printer, the system alerts you with a message. You can determine where the alert message goes (who receives it) when you define the form.

To submit a job for a particular form, include the **-f** option on the **lp** command line. For example, if you have a file called **PO-3992** that is formatted to be printed as a purchase order using a form you call **porder**, print the file on the form with a command line like this:

```
% lp  -f  porder  PO-3992  ↵
```

The Forms menu offers several operations for managing forms:

`Create`	Create a new form definition.
`Modify`	Change an existing form definition.

`Remove`	Delete a form description from the LP services databases.
`Mount`	Assign a form restriction to a printer.
`Unmount`	Remove a form restriction from a printer.
`List`	Display information about the forms currently defined on the system.
`Show Status`	Display the current status of forms on the system.

The following sections elaborate.

Creating Forms

Select the Create operation to enter a forms description. The operation asks a number of questions having to do with the desired page layout and other printing attributes. Some of these attributes, characters per line, for example, depend on the capabilities of your printer. You may need to consult your printer documentation to see what it can do and what it cannot do.

The following sections discuss the information you may specify when you define the form. The default form contains values for printing a normal page of 66x80 ASCII characters.

Name

The form name is any arbitrary name that you may choose. You should select a name that describes the form sufficiently but is still easy to type.

Form to Copy

When defining a form, you may specify an existing form whose definition you wish to start with as a model. If you need to define numerous forms that are similar to each other but that do not use many of the **sysadm**'s preset form defaults, you can create your own default form and copy the others from it.

Page Length

This value specifies the length of the form. For multi-page forms, this value is the length of each page. Specify the length as a number of lines, or in inches or centimeters. To specify a length in inches or centimeters, follow the value with an **i** or **c**, as appropriate. For example, specify six inches as **6i**.

Page Width

This value specifies the width of the form. Specify the length as a number of characters, or in inches or centimeters. To specify a length in inches or centimeters, follow the value with an **i** or **c**, as appropriate. For example, specify 20 centimeters as **20c**.

Number of Pages

This value specifies the number of pages making up a multi-page form.

Line Pitch

This value determines how many lines appear per vertical inch. You may specify this value in either lines per inch or lines per centimeter by following the value with an **i** or a **c**, respectively. If you omit the letter, the operation assumes that you mean inches.

Character pitch

This value specifies the number of characters printed per horizontal inch. You may specify this value in either characters per inch or characters per centimeter by following the value with an **i** or a **c**, respectively. For example, specify 3 characters per centimeter with **3c**. If you omit the letter, the operation assumes that you mean inches.

Print Wheel

For this field, you may make up a name to represent a particular print wheel, character set, or font cartridge required to print the form. Make the name descriptive but easy to type. You should be consistent in your naming conventions.

Users can submit print requests specifying a particular print wheel or character set by including the **–S** option on the **lp** command line. If the required print wheel or character set is not mounted on the destination printer, the LP service will hold the job and alert you that you need to mount the print wheel or character set.

Ribbon Color

If the form requires a particular color of ribbon, you specify it in this field. Whenever you mount the form, the LP service will remind you to load the correct color of ribbon. The LP service does not track ribbons; therefore, it will not alert you if the wrong ribbon is loaded on a printer. It is up to you make sure the correct ribbon is loaded.

Comment

The comment field is for you to enter any information you wish about the form. Users can display the comment for a form; therefore, it is useful for you to enter a comment describing the form, its purpose, and so on.

Alignment Pattern

The alignment pattern is any pattern of characters that you may print to see if the form stock is loaded correctly in the printer. For security reasons, the LP service only allows the superuser and itself to see the alignment pattern.

Alert Messages

When a user submits a request requiring a form that is not mounted on a printer, the LP service alerts you. The alert may occur several different ways:

`mail` Using the **mail**(1) command, the LP service sends mail to the administrator.

`write` Using the **write**(1) command, the LP service writes a message to the terminal where the administrator is currently logged in.

`quiet` Send no alert for the current form need. This option is useful for silencing alerts for a form need that you already know exists. Selecting `quiet` does not change the current alert status if there is currently no form need.

`none` The LP service sends no alerts if it detects a need for a form.

As an alternative, you may write your own alert script. The script should accept the alert message as standard input from the LP service. When prompted, specify the command line for invoking your custom alert script.

You may also specify how often the LP service alerts you when a form need arises. You specify the interval in minutes. For example, an interval of **10** means that the LP service will send out an alert every 10 minutes for as long as the need exists. If you want the LP service to send out only one alert message upon detecting a need for a form, specify **0**.

Job Threshold

If you do not want to receive an alert message every time someone needs a form mounted, you can set a threshold determining how many requests needing a form will occur before the LP service alerts you. For example, setting the threshold to **5** means that the LP service will not alert you to mount a form until it has received 5 print requests needing a form.

Users to allow

Optionally, you may specify which users may use the form Regardless of any users expressly denied access to the form (via `Users to deny`, discussed below), only the users specified in `Users to allow` may use the form. Specifying **all** allows access for all users except any who are denied access in **Users to deny**.

Users to deny

Optionally, you may specify which users may not use the form. If you specify users in `Users to allow`, discussed above, it is not necessary to specify users in `Users to deny`.

Modifying Forms

Use the Modify operation to change the description for an existing form. For a discussion of the fields in the form description, see "Creating Forms."

Removing Forms

To remove the description of a form from your system, select the Remove operation.

Mounting and Unmounting Forms

The LP service tracks forms status by allowing you to mount and unmount forms whenever you load or unload, respectively, a form on a printer. After loading form stock on a printer device, select the Mount operation to tell the LP service and users that the form is now available. With the form mounted, the LP service will print requests requiring the form.

After unloading form stock from a printer, use the Unmount operation to tell the LP service and users that the form is no longer available. If a user submits a request that requires a form that is not mounted, the LP service holds the job in the queue without printing it, and it sends you an alert message. The nature of the alert message depends on the Alert Messages option that you selected when you added the form. For more information on these options, see "Creating Forms," earlier in the chapter. To change the way the LP service currently handles alert messages, select the Modify operation.

To see what forms are mounted on what printers on your system, select the Show Status operation.

Displaying Forms

The List operation displays the form description attributes that you set with the Create operation. These attributes are:

`Name`	The name that you gave the form. This line in the display also tells whether or not the form is available for you to use.
`Page Length`	The length of the page in lines, inches, or centimeters. Values in inches and centimeters are followed by **i** or **c**, respectively.
`Page Width`	The width of the page in characters, inches, or centimeters. Values in inches and centimeters are followed by **i** or **c**, respectively.
`Number of Pages`	The number of pages making up the form.
`Line Pitch`	The number of lines per vertical inch.
`Character Pitch`	The number of characters per horizontal inch.
`Character Set`	The character set, print wheel, or font cartridge required by the form.
`Ribbon Color`	The color of ribbon required for the form.
`Comment`	A comment that you supplied when you added the form. This comment may describe the form, its purpose, and so on.
`Alignment`	A pattern of characters that you can print to see if the form stock is loaded correctly in the printer. For security reasons, only the superuser and the LP service can access and display the alignment pattern.

For a more detailed discussion of these attributes, see "Creating Forms."

Showing Status of Forms

Use the Show Status operation to see what forms on your system are currently mounted and available to you.

Setting User and Request Priorities

The LP service offers several different ways of controlling which jobs print first when multiple jobs are competing for printing resources. This priority mechanism

revolves around the concept of a priority level associated with each job. Very simply, the lower the priority level number, the higher the priority of the request. Conversely, requests with higher priority numbers have lower priority. The LP service prints jobs with lower priority level numbers first.

Using the **–q** option of the **lp** command, users can set priority numbers for the jobs they submit. Priorities range from 0 (high priority) to 39 (low priority). For example, to submit a request to print file **nicefile** at the lowest priority, use this command:

```
% lp -q 39 nicefile ↵
```

To help regulate the use of the priority system, the LP service allows you to set limits for priorities that users can specify. The Priorities menu provides these operations:

`Job Default`	Set a default priority for requests that do not specify a priority with the **–q** option.
`Remove`	Remove a default job priority limit for a specific user.
`Set`	Set a default job priority limit for a specific user.
`User Default`	Set a default job priority limit for all users who do not have an individual limit.
`List`	List current job priority limits for general requests and for specific users.

The following sections describe these operations in more detail.

Job Default

Use the Job Default operation to set the priority for all requests submitted without the **–q** option. You use the **–q** option on the **lp** command line to set the priority for a job.

Remove

Use this operation to remove the priority limit set for an individual user. You set an individual limit with the Set operation. Once you have removed the individual limit, the user's requests are subject to the User Default, if set.

Set

This operation sets a priority limit for an individual user. The lower the priority level, the higher priority job they may submit. Regardless of a user's priority limit, if the user does not specify a priority with the **–q** option, the request takes the system default priority level.

User Default

Use this operation to set a priority limit for all users for whom you have not set an individual limit. Regardless of the user default priority limit, if a user does not

specify a priority with the **–q** option, the request takes the system default priority level.

List

This operation displays the priority values set with the other Priorities menu operations.

Managing Requests

The Requests menu offers several operations for handling jobs in the print queue. Unlike other printer management operations (other than list operations), some of these operations are accessible to users other than the superuser. The following sections elaborate.

Canceling a Request

Use the Cancel operation to remove a request from a queue. The superuser can cancel any request, but other users can cancel only their own requests.

Holding a Request

Any user can use the Hold operation to suspend his or her own request. A held request remains in the queue, but the LP service will not send it to a printer until the user (or superuser) releases the request with the Resume operation. A held request does not block the print queue: the LP service will continue to serve other requests.

Resuming a Held Request

Use the Resume operation to release a request that a user suspended with the Hold operation. Users other than the superuser can resume only their own requests.

Moving Requests

The superuser can use the Move operation to move requests from the queue for one printer or class to the queue for another printer or class.

Displaying Requests

Use the List operation to display the requests currently in the print queue. Like other List operations in **sysadm**, any user may use this operation.

Managing the Scheduler

The LP scheduler is the process that manages print queues and the other LP services. The system starts the LP scheduler when you bring the system up to run level 1. The Scheduler menu offers three operations, described in the next three subsections.

Starting the Scheduler

Use the Start operation to begin execution of the LP scheduler. This operation performs no function if the scheduler is already running.

Stopping the Scheduler

Use the Stop operation to halt execution of the LP scheduler. This operation performs no function if the scheduler is already stopped.

This operation does not remove jobs from the print queue. When you restart the scheduler, it resumes processing queued requests.

Displaying Scheduler Status

Use the List operation to display the status of the LP scheduler and the print queues. The operation tells whether or not the scheduler is running, and it displays the queue for each printer on the system. In a queue display, the entry for each request shows:

- The request number.
- The user who submitted the request.
- The size of the file in bytes.
- The time and date when the job was submitted.
- Any printer specified.

Managing Remote Systems

If your system is on a network, the LP service allows you to share printer services with other systems in your network. Use the operations in the Systems menu to manage the LP service databases that contain information about remote systems.

Before adding a remote printer that resides on a system using a DG/UX Release 5.4 LP scheduler, add an entry for the remote system using the Add operation in the Systems menu. On the remote system, the system administrator must do the same, adding similar information about your system. After adding this information on both systems, you can set up access to the remote system's printers, and the remote system administrator can set up access to your printers.

The following sections discuss the operations in the Systems menu.

Adding Remote Systems

Use this operation to add the required LP service database entries that allow you to share printing resources with a remote system. To add a remote system, you need to do several things:

`Name` See the administrator of the remote system to verify the system's host name. If your systems are not part of an NIS (Network Information Services)

domain, you need to make sure that entries for the host exist in your system's **/etc/ethers** and **/etc/hosts** files. To add entries to these files, see the **sysadm** utility's Networking menu.

If you are in an NIS domain, the required entries may already exist in the NIS databases. If not, see your NIS administrator.

`Scheduler Type`

You need to know if the remote system's scheduler is compatible with the AT&T System V scheduler (as is the DG/UX Release 5.4 scheduler) or with the BSD scheduler. If you are not sure, try the AT&T scheduler (**s5**) first.

`Connection Timeout Period`

Set a time in minutes that the connection with the remote system may remain idle before "timing out," or terminating.

`Connection Retry Interval`

Set the number of minutes to wait, after unexpected disconnection of service from the remote system, before attempting to reconnect.

`Comment`

Decide on an optional comment for the system entry. The List operation will display this comment with the other system information.

Modifying Remote Systems

Use the Modify operation to change the attributes associated with a system entry. See "Adding Systems" for information about the entry attributes.

Removing Remote Systems

The Remove operation deletes the remote system entry from the LP services databases. This operation does not remove entries from the **/etc/ethers** or **/etc/hosts** files.

Displaying Remote Systems

Use the List operation of the Systems menu to display remote system entries added with the Add operation. The section on the Add operation, above, describes the fields in the entry.

Displaying LP Service Status

The List operation of the Printer menu displays the status of the entire LP services subsystem. The display includes:

- Whether the scheduler is running.
- Which printer, if any, is the default destination.
- The device/printer name assignments.
- The accept status of each printer and the time and date when the accept status last changed.

- Whether the printer is enabled or disabled and the time and date when the enabled/disabled status last changed. If disabled, the report includes the reason.
- Which forms, if any, are available to you, and where they are mounted.
- Requests currently in the queue.

Expert Information

This section describes the shell-level commands that **sysadm** uses to provide LP services. Table 12–1 shows which **sysadm** menus and shell commands are available for administering the LP service.

Table 12–1 LP Print Service Menu and Command Summary

Task Description	Menu Selection	Shell Command
Configure printers for print service, set the default printer, turn queuing of requests on and off, enable and disable printers, display state and configuration of printers	**Devices**	**lpadmin**(1M), **accept**(1M), **reject**(1M), **enable**(1), **disable**(1), **lpstat**(1)
Group printers into classes	**Classes**	**lpadmin**(1M)
Provide pre-processing software for files to be printed	**Filters**	**lpfilter**(1M)
Define pre-printed forms for print requests	**Forms**	**lpforms**(1M)
Define levels of priority available to users requesting print jobs	**Priorities**	**lpusers**(1M)
Cancel, hold, resume, move, and list print requests	**Requests**	**lp**(1), **cancel**(1), **lpmove**(1M), **lpsched**(1M)
Start and stop print service, report status of scheduler and list print requests	**Scheduler/Service**	**lpsched**(1M), **lpshut**(1M), **lpstat**(1)
Set up communication to remote print service	**Systems**	**lpsystem**(1M)
Identify active printers, print wheels & character sets, mounted forms, and pending requests	**List**	**lpstat**(1)

The rest of this section describes the work required to set up and maintain print services with the LP print service utilities. Details about the commands listed above are available in the manual pages for them.

This section includes the following information:

- A description of how the LP print service works
- References to documentation for installing the print service
- Troubleshooting guidelines
- Instructions for stopping and starting the print service manually
- Instructions for configuring a print service for the unique requirements of your users (such as the need for particular pre-printed forms and filters)
- A list of directories and files delivered as part of the print service
- Instructions for supporting PostScript printers
- Instructions for writing customized filters and interface programs

Overview

The LP print service, originally called the LP spooler, is a set of software utilities that allows you, minimally, to send a file to be printed while you continue with other work. The term "spool" is an acronym for "simultaneous peripheral output on-line," and "LP" originally stood for Line Printer, but has come to include many other types of printing devices. The print service has many optional enhancements; however, you can make your print service as simple or as sophisticated as you like.

Components of a Print Service

A print service consists of both hardware and software. You must have at least one computer and one printing device for an LP print service. Beyond that, you may have any number of computers and printing devices; there is no limit to the number of pieces of hardware you may include. The software consists of the LP print service utilities and any filters (programs that process the data in a file before it is printed) that you may provide. Users of your print service may be required to print all their files in the same format, or, if you make different types of printers and/or filters available with your service, they may choose from several formats. You may also offer your users a choice between plain paper and pre-printed forms (such as invoices or checks).

Functions Performed by the Print Service Software

Whether your print service is simple (such as a one-computer/one-printer configuration that prints every file in the same format on the same type of paper) or a sophisticated one (such as a computer network with multiple printers and a choice of printing formats and forms), the LP software helps you maintain it by performing several important functions:

- Scheduling the print requests of multiple users
- Scheduling the work of multiple printers

- Starting programs that interface with the printers
- Filtering users' files (if necessary) so they will be printed properly
- Keeping track of the status of jobs
- Keeping track of forms and print wheels currently mounted and alerting you to mount needed forms and print wheels
- Alerting you to printer problems

Suggestions for LP Print Service Administration

Here are some tips about how to organize your duties as the administrator of an LP print service.

Configuring Printer Sites

Where you decide to put your printers and how you decide to connect them to your computers depends on how those printers will be used. There are three possible scenarios: (1) users may access printers attached to their own computer; (2) users may access printers attached to a server computer; and (3) users may access remote printers on a network to which their computer belongs.

- You may want to connect a particular printer directly to the computer that is the home system of the users who will use that printer most often. If you do, the type of connection you have will be referred to as a direct connection. An environment that includes more than one computer, each of which has a direct connection to a printer, is said to have a "distributed printing configuration."
- You may want to have all your printers in one physical location, such as a computer center. If so, you might connect them all to one computer. Users on other computers who want to use a printer may access it through a network linking their own computers to the computer serving the printers. An environment in which one computer serves several printers (which can be accessed only through a computer-to-computer network) is described as a "print server configuration."

Figure 12–2 shows a sample print server configuration.

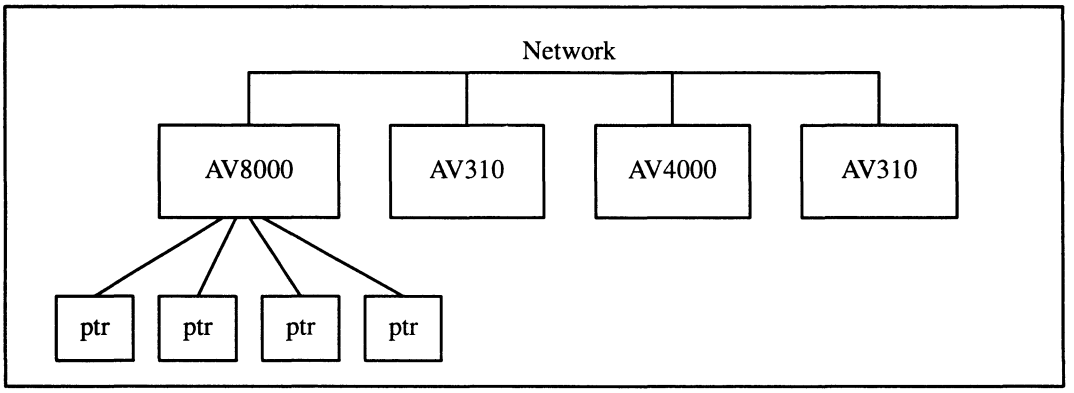

Figure 12–2 Print Server Configuration

You may want to link most of your printers to a dedicated printer server computer, while allowing other printers to be connected to your system. If so, you can arrange your computers and printers in a network configuration as shown in Figure 12–3.

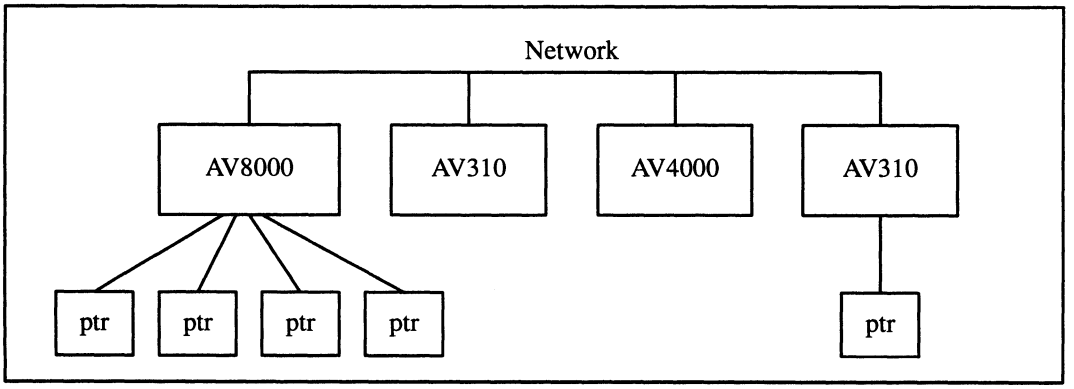

Figure 12–3 Network Configuration

Configuring Printers

Before the LP print service can start accepting print requests, you will have to describe the characteristics of each printer you have. The following is a list of the attributes most commonly defined:

- printer name (mandatory)
- connection method (mandatory for local printers)
- system name (mandatory for access to remote printers and mandatory for allowing remote access to local printers)
- interface program
- printer type
- content types
- printer port characteristics
- character sets or print wheels
- alerting to mount a print wheel
- forms allowed
- printer fault alerting
- printer fault recovery
- restrictions on user access
- inclusion of banner page in output

- printer description
- default printing attributes
- printer class membership
- system default destination
- mounting a form or print wheel
- removing a printer or class

You need to specify very little of this information to add a new printer to the LP print service. The more information you provide, however, the better the printer will satisfy various users' needs.

The descriptions in the sections below will help you understand what this printer configuration information means and how it is used, so that you can decide how to configure your printers. In each section you will also be shown how to specify this information when adding a printer. While you can follow each of the sections in order and correctly configure a printer in several steps, you may want to wait until you've read all the sections before adding a printer, so that you can do it in fewer steps.

Printer Name

The printer name and the connection method (described next) are the only items you must specify to define a new local printer. To define a new remote printer, you must specify the printer name and the name of the remote system. The printer name is used to identify the printer, both by you (when you want to change the printer configuration or manage the printer), and by users who want to use the printer. The name may contain a maximum of 14 alphanumeric characters and underscores.

You may choose any name you like, but it is good practice to choose a name that is meaningful to the users of the LP print service. For example, **laser** is a good name for a laser printer. If you have several laser printers you may name them **laser1**, **laser2**, and so on.

You don't have to try to fit a lot of descriptive information into the name; there is a better place for this information (see "Printer Description" below). You also don't have to make the name precisely identify the type of printer; users who need to use a particular type of printer can specify it by type rather than name (see "Printer Type," below).

You will use the printer name every time you want to refer to the printer: when adding other configuration information for the printer, when changing the configuration of the printer, when referring to the status of the printer, and when removing the printer. Thus the first thing you must do to add a printer is identify its name. You will do this as shown below; but don't do it yet because you'll also need to specify the connection method.

```
# lpadmin -p printer-name ↵
```

There are no default names; you must name every printer.

Connection Method

This section does not apply if you are making a remote printer (one that is connected to a remote system on your network) accessible to users on your system. The LP print service allows you to connect a printer to your computer in one of the following three ways:

- by connecting the printer directly to your computer
- by connecting the printer directly to a network to which your computer is attached
- by connecting the printer to a modem.

Figure 12–4 shows these three types of connections.

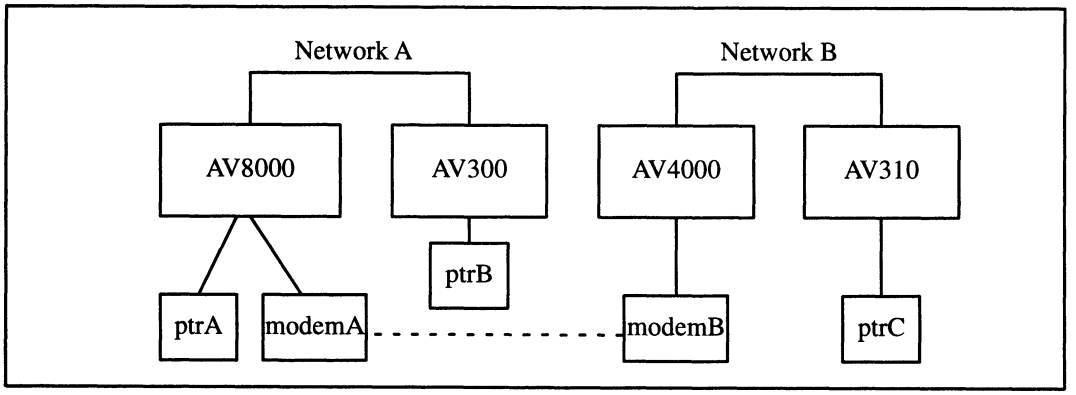

Figure 12–4 Methods of Connecting a Printer to a Computer

The AViiON 8000 system accesses printer **ptrA** through a direct connection and accesses printer **ptrB** over the network. The AViiON 300 system access printer **ptrB** via direct connection and printer **ptrA** over the network. The AViiON 4000 and AViiON 310 systems can both access printer **ptrC**, the former via the network and the latter via direct connection. The AViiON 8000 and AViiON 4000 systems have modems, meanwhile, which allows either system on either network to access printers on the other network. Modem connections are discussed in the next section.

The simplest connection method is by connecting a printer directly to your computer. You may, however, want to connect a printer to a network (so it can be shared with other computers or workstations), or to a modem. Whichever method you use, you must describe it to the LP print service after you have connected the hardware.

To define the connection method for a new printer for your print service, invoke the **lpadmin** command, specifying a connection method through either the **–v** option for a directly connected printer or the **–U** option for a printer directly connected to a network or a printer connected to a modem.

Direct Connections

The simplest and most common method by which printers are connected to a computer is direct connection. If you use this method, you generally need to specify

only two items on the command line when you make the connection: the name of the printer and the name of the connecting port. To connect a printer directly to your computer enter the following command:

```
# lpadmin -p printer-name -v pathname ↵
```

where *pathname* is the name of the special device file representing the printer port. The following are examples of typical names of special device files.

```
/dev/tty00  (serial)
/dev/tty16  (serial)
/dev/lp  (parallel)
```

(For details about using these files, see "Printer Port Characteristics.")

Using a Printer As a Login Terminal

Some directly connected printers can also be used as terminals for login sessions. If you want to use a printer as a terminal, you must arrange for the LP print service to handle it as such. To do so, use the **–l** option to the **lpadmin** command, as follows:

```
# lpadmin -p printer-name -v pathname -l ↵
```

As before, *pathname* is the name of the special file representing the printer port. If the **–l** option is specified, the printer will be disabled automatically whenever the LP print service is started, and therefore will have to be manually enabled before it can be used for printing. For instructions on manually enabling a printer, see "Enabling and Disabling a Printer" (under "Making Printers Available") later in this section.

Connections to Networks and Modems

In an environment where a printer is located is so far from the computer that a direct connection is not possible or practical, you can use a modem or network to access the printer. For example, you might have one printer in use with a single terminal at a branch office located a few miles from your main site, or you may want to share a printer with computers that are not on a common network.

The LP print service establishes a connection to the printer as necessary to print requests; at the end of each request the connection is dropped, making the printer available to the next system that calls it. Thus the printer gets shared by the users of all the computers, more or less equally.

There are two methods for connecting printers that are not directly connected to your system: attached directly to a network and through a dial-up modem. The LP print service uses UUCP to handle both methods.

When a modem connection is used, the printer must be connected to a dialed modem, and the dial-out modem must be connected to the computer. Whether printers are connected to a modem or directly to a network, the connection must be described to UUCP. For instructions on describing either type of connection, see the chapter on network management.

To make a printer connected in one of these ways available to your users, enter the following command:

```
# lpadmin -p printer-name -U dial-info ↵
```

where *dial–info* is either the telephone number to be dialed to reach the printer's modem or the system name entered in the UUCP **Systems** database for the printer. The **–U** option provides a way to link a single printer to your print service. It does not allow you to connect with a print service on another system.

A note on printers connected to a modem or directly to a network: if the printer or port is busy, the LP print service will automatically retry later. This retry rate is 10 minutes if the printer is busy, and 20 minutes if the port is busy. These rates are not adjustable; however, you can force an immediate retry by issuing the **enable** command for the printer. If the port or printer is likely to be busy for an extended period, you should issue the **disable** command.

The **lpstat –p** command reports the reason for a failed dial attempt. Also, if you are alerted to a dialing fault (see "Printer Fault Alerting," later in this section), the alert message will give the reason for the fault. These messages are identical to the error messages produced by UUCP for similar problems. See Appendix A for a list of UUCP error messages and explanations.

In summary, to add printers to your system, use the **lpadmin** command, specifying a connection method through one of two options: the **–v** option for a directly connected printer or the **–U** option for a networked printer or a modem–connected printer.

System Name

This section does not apply if you are making only a local printer accessible to users on your system. A remote printer is one that is connected to a system other than your local system that you can access only through that remote system.

There are some exceptional cases, however. For example, if only one of the printers has a particular typesetter needed for some print jobs, then users from many systems will want to access it from time to time.

Alternatively, a large community of users on a local area network may want to pool all printers on a single system, where they can share them. When this is done, the system supporting the printers becomes a printer server.

To make accessible a printer that is remote, the name of the system on which the printer resides must be registered with the print service. If the remote printer resides on a DG/UX Release 5.4 system or on a System V Release 4 system, enter:

```
# lpsystem system-name ↵
```

If the remote printer resides on a BSD system, enter

```
# lpsystem -t bsd system-name ↵
```

NOTE: If your system is using fully qualified names, then the full name must be used in the command line. If your network contains hosts using both simple and fully qualified names, then both the simple and full name should be added for each print server and client.

For details about the options available with this command, see the **lpsystem**(1M) manual page.

In either case, after entering the **lpsystem** command, enter the **lpadmin** command, as follows:

```
# lpadmin -p printer -s system-name ↵
```

where *printer* is the name by which your users identify the remote printer. You can usually use the same name used for that printer by the remote system. If the name used by the remote system is the same name used for an existing printer or class on your system, you must use a different name. To assign a different name to a remote printer, enter the following:

```
# lpadmin -p local-name -s system-name!remote-name ↵
```

For example, imagine you want your users to have access to a printer called **psjet2** that resides on a remote system called **newyork**. Because you already have a printer called **psjet2** on your own system, you want to give the remote printer a new name on your system: **psjet3**. Request the new name by entering the following:

```
# lpadmin -p psjet3 -s newyork!psjet2 ↵
```

Before you add a remote printer to your system, be sure communications between your system and the network have been properly set up and verified.

Allowing Remote Users to Access Local Printers

Making the printers on your local system accessible to users on remote systems is a two-step process: you must configure the port monitor on the local system, and you must register the remote system with the local LP print service. This section provides instructions for these tasks.

Configuring the Local Port Monitor

If a remote system requires access to printers connected directly to your system, you need to configure the local port monitor for the network you share to accept service requests and to notify the LP print service of such requests. For DG/UX Release 5.4 systems or System V systems calling your systems, issue the following command:

```
# pmadm -a -p netname -s lp -i root -v 'nlsadmin -V' -m 'nlsadmin \
    -o /var/spool/lp/fifos/listenS5'  ↵
```

where *netname* is the name of a network such as **tcp**.

If you expect users on BSD systems to send print requests to your system, then you need to configure your local port monitor. First, however, you need to know the Internet address of your system. To get this address in the correct hexadecimal format, run the **lpsystem** command with the **-A** option, as follows:

```
# lpsystem -A ↵
```

You then execute the following command, substituting the hexadecimal number for the argument address:

```
# pmadm -a -p tcp -s lpd -i root -v 'nlsadmin -V' -m 'nlsadmin \
     -o /var/spool/lp/fifos/listenBSD -A'xaddress'' ↵
```

Adding a System Entry

If you want your system to accept jobs from a remote system (and vice-versa), the print service must know about that system. The **lpsystem** command allows you to register remote systems with the local print service. Run the command as follows:

```
# lpsystem system-name ↵
```

NOTE: If your system is using fully qualified names, then the full name must be used in the command line. If your network contains hosts using both simple and fully qualified names, then both the simple and full name should be added for each print server and client.

where *system–name* is the name of the remote system.

If the remote host uses a pre-DG/UX Release 5.4 LP scheduler, you also need to add the host's Internet address to your **hosts** database. To add an entry to your **hosts** database, use the TCP/IP Databases menu of **sysadm**'s Networking menu.

Interface Program

This section does not apply if you are making a remote printer accessible to users on your system. This is the program the LP print service uses to manage the printer each time a file is printed. It has several tasks:

- to initialize the printer port (the connection between the computer and the printer)
- to initialize the printer (restore it to a normal state in case a previously printed file has left it in an unusual state) and set the character pitch, line pitch, page size, and character set requested by the user
- to print a banner page
- to run a filter that prepares the file for printing
- to manage printer faults

If you do not choose an interface program, the standard one provided with the LP print service will be used. This should be sufficient for most of your printing needs. If you prefer, however, you can change it to suit your needs, or completely rewrite your own interface program, and then specify it when you add a new printer. See "Customizing the Print Service" later in this section for details on how to customize an interface program.

If you are using the **standard** interface program, you needn't specify it when adding a printer. If, however, you will be using a different interface program on a

local printer, you can refer to it either by specifying its full pathname or by referring to another printer using the same interface program.

To identify a customized interface program by name, specify the printer name and the pathname of the interface program, as follows:

```
# lpadmin -p printer-name -i pathname ↵
```

To use a customized interface program of another printer, specify the printer names as follows:

```
# lpadmin -p printer-name1 -e printer-name2 ↵
```

Printer–name1 is the name of the printer you are adding; *printer–name2* is the name of an existing printer that is using the customized interface program.

Printer Type

A printer type is the generic name for a printer. When you set up your system you can enhance its ability to serve your users by classifying, on the basis of type, the printers available through the print service. Assigning a type for each printer is also important because the LP software extracts information about printers from the **terminfo** database on the basis of type. This information includes the list of the printer's capabilities that is used to check the configuration information you supply to the print service. (By checking information provided by you against the capabilities of the printer, the print service can catch any inappropriate information you may have supplied.) The **terminfo** database also specifies the control data needed to initialize a particular printer before printing a file.

You can assign several types to a printer if the printer is capable of emulating more than one kind of printer. If you specify more than one printer type, the LP print service will use one of them as appropriate for each print request.

Although you do not need to specify a printer type, we recommend that you do so; when a printer type is specified, better print services can be provided.

To specify a printer type, use the following command line:

```
# lpadmin -p printer-name -T printer-type-list ↵
```

If you give a list of printer types, separate the names with commas. If you do not define a printer type, the default **unknown** will be used.

Content Types

While the printer type tells the LP print service what types of printers are being added, the content types tell the LP print service what types of files can be printed directly on each printer. Most printers can print files of two types: the same type as the printer type (if the printer type is defined) and the type **simple**, (meaning an ASCII file) which is the default content type for all printers.

Some printers, though, can accept (and print properly) several different types of files. When adding this kind of printer, specify the names of the content types the new printer accepts by adding these names to the list. (By default, the list contains only one type: **simple**.) If you are adding a remote printer, enter the following command on the remote system to list the content types that it accepts:

```
# lpstat -p printer -l
```

To specify the list of content types for the local printer, enter the following command:

```
# lpadmin -p printer-name -I content-type-list
```

The *content–type–list* is a list of names separated by commas or spaces. If you use spaces to separate the names, enclose the entire list (but not the **–I**) in quotes.

Content type names may look a lot like printer type names, but you are free to choose names that are meaningful to you and the people using the printer. (The names **simple** and **any** are recognized as having particular meanings by the LP print service; be sure to use them consistently. The name **terminfo** is also reserved, as a reference to *all* types of printers.) The names must contain no more than fourteen characters and may include only letters, digits, and underscores. Table 12–2 describes some accepted content types.

Table 12–2 Content Types

Types	Description
troff	Device independent output from **troff**
otroff	CAT typesetter instructions generated by BSD (old) **troff**
tex	DVI format files
plot	Plotting instructions for Tektronix displays and devices
raster	Raster bitmap format for Varian raster devices
cif	Output of BSD **cifpbt**
fortran	ASA carriage control format
epsonfx	Epson and compatible printers
proprinter	IBM ProPrinter and compatible printers
laserjet	Hewlett-Packard LaserJet and compatible printers
PS	PostScript language
simple	ASCII file

When a file is submitted to the LP print print service for printing with the printer specified by the **–d any** option of the **lp** command, the print service searches for a

printer capable of handling the job. The print service can identify an appropriate printer through either the content type name or the printer type name. Therefore, you may specify either name (or no name) when submitting a file for printing. If the same content type is printable by several different types of printers, you should use the same content type names when you add those printers. This makes it easier for the people using the printers, because they can use the same name to identify the type of file they want printed regardless of the printing destination.

Most manufacturers produce printers that accept simple ASCII files. While these printers are different types (and thus have different initialization control sequences), they may all be capable of handling the same type of file, which we call **simple**. As another example, several manufacturers may produce printers that accept ANSI X3.64 defined escape sequences; however, the printers may not support all the ANSI capabilities; they may support different sets of capabilities. You may want to differentiate them by assigning different content types to these printers.

However, while it may be desirable (in situations such as these) to list content types for each printer, it is not always necessary to do so. If you don't, the printer type will be used as the name of the content type the printer can handle. If you have not specified a printer type, the LP print service will assume the printer can print only files of content type **simple**. This may be sufficient if you require users to specify the proper printer explicitly and if files are properly prepared for the printer before being submitted for printing.

The Default Content Type: simple

Files of content type simple are assumed to contain only two types of characters, printable ASCII characters and the following control characters:

backspace	moves the carriage back one space, except at the beginning of a line.
tab	moves the carriage to the next tab stop; by default, stops are spaced every 8 columns on most printers.
linefeed	moves the carriage to the beginning of the next line (may require special port settings for some printers—see "Printer Port Characteristics" below).
form feed	moves the carriage to the beginning of the next page.
carriage return	moves the carriage to the beginning of the same line (may fail on some printers).

The word "carriage" may be archaic for modern laser printers, but these printers do actions similar to those done by a carriage. If a printer can handle several types of files, including **simple**, you must include simple in the content type list; the type **simple** is not automatically added to any list you give. If you *don't* want a printer to accept files of type **simple**, give a blank *content–type–list*, as follows:

```
# lpadmin -p printer-name -I ""  ↵
```

Printer Port Characteristics

This section applies only to local printers.

The interface program needs to set the characteristics of the port to which the printer is connected. These characteristics define the low level communications with

the printer. Included are the baud rate; use of XON/XOFF flow control; 7, 8, or other bits per byte; type of parity; and output post-processing. The standard interface program will use the **stty** command to initialize the printer port, minimally setting the baud rate and a few other default characteristics.

The default characteristics applied by the standard interface program are listed below.

Default	Meaning
9600	9600 baud rate
cs8	8-bit bytes
–cstopb	1 stop bit per byte
–parenb	no parity generation
ixon	enable XON/XOFF flow control
–ixany	allow only XON to restart output
opost	post-process data stream as listed below:

–olcuc	don't map lower-case to upper-case
onlcr	map linefeed into carriage-return/linefeed
–ocrnl	don't map carriage-return into linefeed
–onocr	output carriage-returns even at column 0
nl0	no delay after linefeed
cr0	no delay after carriage-return
tab0	no delay after tab
bs0	no delay after backspace
vt0	no delay after vertical tab
ff0	no delay after form-feed

NOTE: If you are printing binary or graphics files, use the **stty** option **–opost** rather than the default **opost**.

You may find that the default characteristics are sufficient for your printers; however, printers vary enough that you are likely to find that you have to set different characteristics. For a complete list of characteristics, see the **stty**(1) manual page.

If you have a printer that requires printer port characteristics other than those handled by the **stty** program, you may have to customize the interface program. See "Customizing the Print Service" for help. You can set the characters for a specific print request using the **lp** command option **–o stty=***stty-option-list*. You can also use

the **lpadmin** command option **–o stty=***stty-option-list* to set the characteristics on a per print queue basis.

When you add a new printer, you may specify an additional list of port characteristics. The list you give will be applied after the default list, so that you do not need to include in your list items that you don't want to change. Specify the additional list as follows:

```
# lpadmin -p printer-name -o "stty='stty-option-list'" ↵
```

Note that both the double quotes and single quotes are needed if you give more than one item in the *stty–option–list.*

As one example, suppose your printer is to be used for printing graphical data, where linefeed characters should be output alone, without an added carriage-return. You would enter the following command:

```
# lpadmin -p printer-name -o "stty=-onlcr" ↵
```

Note that the single quotes are omitted because there's only one item in the list.

As another example, suppose your printer requires odd parity for data sent to it. You would enter the following command:

```
# lpadmin -p printer-name -o "stty='parenb parodd cs7'" ↵
```

Character Sets or Print Wheels

Although your users may use character sets or print wheels that have been mounted on a remote printer (by the administrator of the remote system on which that printer resides), you cannot mount a character set or a print wheel on a remote printer. Printers differ in the way they can print in different font styles. Some have changeable print wheels, some have changeable font cartridges, others have preprogrammed, selectable character sets.

When adding a printer, you may specify what print wheels, font cartridges, or character sets are available with the printer. If you're adding a remote printer and you want your users to be able to use character sets or print wheels that have been mounted by the administrator of the remote system, you must list those character sets and print wheels, just as you would list the character sets and print wheels on a local printer. (See instructions below.)

Only one of these is assumed to apply to each printer. From the point of view of the LP print service, however, print wheels and changeable font cartridges are the same because they require you to intervene and mount a new print wheel or font cartridge. Thus, for ease of discussion, only print wheels and character sets will be mentioned.

When you list the print wheels or character sets available, you will be assigning names to them. These names are for your convenience and the convenience of the users. Because different printers may have similar print wheels or character sets,

you should use common names for all printers. This allows a user to submit a file for printing and ask for a particular font style, without regard for which printer will be used or whether a print wheel or selectable character set is used.

If the printer has mountable print wheels, you need only list their names. If the printer has selectable character sets, you need to list their names and map each one into a name or number that uniquely identifies it in the **terminfo** database. Use the following command to determine the names of the character sets listed in the **terminfo** database.

```
# tput -T printer-type csnm 0 ↵
```

Printer–type is the name of the printer type in question. The name of the 0th character set (the character set obtained by default after the printer is initialized) will be printed. Repeat the command, using **1**, **2**, **3**, and so on in place of the **0**, to see the names of the other character sets. In general, the **terminfo** names should closely match the names used in the user documentation for the printer; however, because not all manufacturers use the same names, the **terminfo** names may differ from one printer type to the next. For the LP print service to be able to find the names in the **terminfo** database, you must specify a printer type. See "Printer Type" above.

To specify a list of print wheel names when adding a printer, enter the following command.

```
# lpadmin -p printer-name -S print-wheel-list ↵
```

The *print–wheel–list* is a comma or space separated list of names. If you use spaces to separate the names, enclose the entire list (but not the **–S**) in quotes.

To specify a list of character set names and to map them into **terminfo** names or numbers, enter the following command:

```
# lpadmin -p printer-name -S character-set-list ↵
```

The *character–set–list* is also a comma or space separated list; however, each item in the list looks like one of the following:

```
csN=character-set-name
character-set-name1=character-set-name2
```

The *N* in the first case is a number from 0 to 63 that identifies the number of the character set in the **terminfo** database. The *character–set–name*1 in the second case identifies the character set by its **terminfo** name. In either case the name to the right of the equal sign (=) is the name you may use as an alias of the character set. You do not have to provide a list of aliases for the character sets if the **terminfo** names are adequate. You may refer to a character set by number, by **terminfo** name, or by your alias.

For example, suppose your printer has two selectable character sets (sets #1 and #2) in addition to the standard character set (set #0). The printer type is **5310**. You enter the following commands to determine the names of the selectable character sets.

```
% tput -T 5310 csnm 1  ↵
english
% tput -T 5310 csnm 2  ↵
finnish
```

The words `english` and `finnish`, the output of the commands, are the names of the selectable character sets. You feel that the name `finnish` is adequate for referring to character set #2, but better names are needed for the standard set (set #0) and set #1. You enter the following command to define synonyms.

```
# lpadmin -p printer-name -S "cs0=american, english=british" ↵
```

The following three commands will then produce identical results.

```
# lp -S cs1 -d any ...  ↵

# lp -S english -d any ...  ↵

# lp -S british -d any ...  ↵
```

If you do not list the print wheels or character sets that can be used with a printer, then the LP print service will assume the following: a printer that takes print wheels has only a single, fixed print wheel, and users may not ask for a special print wheel when using the printer; and a printer that has selectable character sets can take any **cs***N* name or **terminfo** name known for the printer.

Alerting to Mount a Print Wheel

This section does not apply if you are making a remote printer available to users on your system.

If you have printers that can take changeable print wheels, and have listed the print wheels allowed on each, then users will be able to submit a print request to use a particular print wheel. Until it is mounted though (see "Mounting a Form or Print Wheel" in this section), a request for a print wheel will stay queued and will not be printed. You could periodically monitor the number of print requests pending for a particular print wheel, but the LP print service provides an easier way: you can ask to be alerted when the number of requests waiting for a print wheel has exceeded a specified threshold.

You can choose one of several ways to receive an alert.

- You can receive an alert via electronic mail. For more information on the mail system, see the **mailx**(1) manual page.
- You can receive an alert written to any terminal on which you are logged in. See the manual page for the **write**(1) command.
- You can receive an alert through a program of your choice.
- You can receive no alerts.

If you elect to receive no alerts, you are responsible for checking to see whether any print requests haven't printed because the proper print wheel isn't mounted. In addition to the method of alerting, you can also set the number of requests that must be queued before you are alerted, and you can arrange for repeated alerts every few minutes until the print wheel is mounted. You can choose the rate of repeated alerts, or you can opt to receive only one alert for each print wheel.

To arrange for alerting to the need to mount a print wheel or character set, enter one of the following commands:

```
# lpadmin -S print-wheel-names -A mail -Q requests -W minutes ↵
# lpadmin -S print-wheel-names -A write -Q requests -W minutes ↵
# lpadmin -S print-wheel-names -A 'command' -Q requests -W minutes ↵
```

The `print-wheel-names` argument may be a space- or comma-separated list of print wheels or character sets. The first two commands direct the LP print service to send you a mail message or write the message directly to your terminal, respectively, for each alert. The third command directs the LP print service to run the *command* for each alert. The shell environment currently in effect when you enter the third command is saved and restored for the execution of *command*; this includes the environment variables, user and group IDs, and current directory. The argument *requests* is the number of requests that need to be waiting for the print wheel before the alert is triggered, and the argument *minutes* is the number of minutes between repeated alerts.

If you do not want the print service to issue an alert when a print wheel needs to be mounted, enter the following:

```
# lpadmin -S print-wheel-names -A none ↵
```

If you want mail sent or a message written to another user when a printer fault occurs, use the third command with the option **–A'mail** *username*' or **–A'write** *username*'.

When you start receiving repeated alerts, you can direct the LP print service to stop sending you alerts (for the current case only), by executing the following command.

```
# lpadmin -S print-wheel-names -A quiet ↵
```

Once the print wheel has been mounted and unmounted again, alerts will start again if too many requests are waiting. Alerts will also start again if the number of requests waiting falls below the **–Q** threshold and then rises up to the **–Q** threshold again, as when waiting requests are canceled, or if the type of alerting is changed.

If *print–wheel–name* is **all** in any of the commands above, the alerting condition will apply to all print wheels for which an alert has already been defined.

If you don't define an alert method for a print wheel, you will not receive an alert to mount it. If you do define a method without the **–W** option, you will be alerted once for each occasion.

Forms Allowed

For information about how to define, mount, and set up alerting to mount a form, see "Providing Forms" later in this section.

You can control the use of preprinted forms on any printer, including remote printers. (Although you cannot mount forms on remote printers, your users may use forms on remote printers—if the form is defined exactly the same on the local and remote systems.) You may want to do this, for instance, if a printer is not well suited for printing on a particular form because of low print quality, or if the form cannot be lined up properly in a local printer.

The LP print service will use a list of forms allowed or denied on a printer to warn you against mounting a form that is not allowed on the printer. However, you have the final word on this; the LP print service will not reject the mounting. The LP print service will, however, reject a user's request to print a file on a printer using a form not allowed on that printer. If, however, the printer is a local printer and the requested form is already mounted, the request will be printed on that form.

If you try to allow a form for a printer, but the printer does not have sufficient capabilities to handle the form, the command will be rejected.

The method of listing the forms allowed or denied for a printer is similar to the method used to list users allowed or denied access to the **cron** and **at** facilities. See the manual page for **cron**(1M) and **crontab**(1). Briefly, the rules are as follows:

- An allow list is a list of forms that are allowed to be used on the printer. A deny list is a list of forms that are not allowed to be used on the printer.
- If the allow list is not empty, only the forms listed are allowed; the deny list is ignored. If the allow list is empty, the forms listed in the deny list are not allowed. If both lists are empty, there are no restrictions on which forms may be used other than those restrictions that apply to a printer of a particular type, such as a PostScript printer for which a license is required.
- Specifying **all** in the allow list allows all forms; specifying **all** in the deny list denies all forms.

You can add names of forms to either list using one of the following commands:

```
# lpadmin -p printer-name -f allow:form-list ↵
# lpadmin -p printer-name -f deny:form-list ↵
```

The *form–list* is a comma or space separated list of names of forms. If you use spaces to separate names, enclose the entire list (including the **allow:** or **deny:** but not the **–f**) in quotes.

The first command shown above adds names to the allow list and removes them from the deny list. The second command adds names to the deny list and removes them from the allow list. To make the use of all forms permissible, specify **allow:all**; to deny permission for all forms, specify **deny:all**.

If you do not use this option, the LP print service will consider that the printer denies the use of all forms. It will, however, allow you to mount any form, thereby making it implicitly available to use. (See "Mounting a Form or Print Wheel" later in this section for more information.)

Printer Fault Alerting

This section does not apply if you are making a remote printer accessible to users on your system. The LP print service provides a framework for detecting printer faults and alerting you to them. Faults can range from simple problems, such as running out of paper or ribbon, or needing to replace the toner, to more serious faults, such as a local power failure or a printer failure. The range of fault indicators is also broad, ranging from dropping carrier (the signal that indicates that the printer is on line), to sending an XOFF, to sending a message. Only two classes of printer fault indicators are recognized by the LP print service itself: a drop in carrier and an XOFF not followed in reasonable time by an XON. However, you can add filters that recognize any other printer fault indicators, and rely on the LP print service to alert you to a fault when the filter detects it. For a description of how to add a filter, see "Providing Filters" in this section. For a description of how a filter should let the LP print service know a fault has occurred, see "Customizing the Print Service" in this section.

You can choose one of several ways to receive an alert to a printer fault:

- You can receive an alert via electronic mail. See the manual page for the **mailx**(1) command.
- You can receive an alert written to any terminal on which you are logged in. See the manual page for the **write**(1) command.
- You can receive an alert through a program of your choice.
- You can receive no alerts.

If you elect to receive no alerts, you will need a way of finding out about the faults and fixing them; the LP print service will not continue to use a printer that has a fault. In addition to the method of alerting, you can also arrange for repeated alerts every few minutes until the fault is cleared. You can choose the rate of repeated alerts, or you can choose to receive only one alert per fault. Without a filter that provides better fault detection, the LP print service cannot automatically determine when a fault has been cleared except by trying to print another file. It will assume that a fault has been cleared when it is successfully able to print a file. Until that time, if you have asked for only one alert per fault, you will not receive another alert. If, after you have fixed a fault, but before the LP print service has tried printing another file, the printer faults again, or if your attempt to fix the fault fails, you will not be notified. Receiving repeated alerts per fault, or requiring manual re-enabling of the printer (see "Printer Fault Recovery," below), will overcome this problem.

To arrange for alerting to a printer fault, enter one of the following commands:

```
# lpadmin -p printer-name -A mail -W minutes ↵
# lpadmin -p printer-name -A write -W minutes ↵
# lpadmin -p printer-name -A'command' -W minutes ↵
```

The first two commands direct the LP print service to send you a mail message or write the message directly to your terminal, respectively, for each alert. The third command directs the LP print service to run the *command* for each alert. The shell environment currently in effect when you enter the third command is saved and restored for the execution of *command.* The environment includes environment variables, user and group IDs, and current directory. The *minutes* argument is the number of minutes between repeated alerts.

If you do not want the LP print service to issue an alert when a fault occurs, enter the following:

```
# lpadmin -p printer-name -A none ↵
```

If you want mail sent or a message written to another user when a printer fault occurs, use the third command with the option **–A'mail** *username*' or **–A'write** *username*'.

Once a fault occurs and you start receiving repeated alerts, you can direct the LP print service to stop sending you alerts (for the current fault only), by executing the following command:

```
# lpadmin -p printer-name -A quiet ↵
```

If the *printer–name* is **all** in any of the commands above, the alerting condition will apply to all printers.

If you don't define an alert method, you will receive mail once for each printer fault. If you define a method without the **–W** option, you will be alerted once for each fault.

Printer Fault Recovery

This section does not apply if you are making a remote printer accessible to users on your system. When a printer fault has been fixed and the printer is ready for printing again or when the fault is unrecoverable, the LP print service will respond in one of four ways:

- Continue printing at the top of the page where printing stopped.
- Restart printing at the beginning of the print request that was active when the fault occurred.
- Wait for you to tell the LP print service to re-enable the printer.
- Abort the job and not re-schedule it.

The ability to continue printing at the top of the page where printing stopped requires the use of a filter that can wait for a printer fault to be cleared before resuming properly. Such a filter probably must have detailed knowledge of the control sequences used by the printer so it can keep track of page boundaries and know where in a file printing stopped. ***Note that the default filter used by the LP print service cannot do this***. If a proper filter is not being used, you will be notified in an alert if recovery cannot proceed as you want.

To specify the way the LP print service should recover after a fault has been cleared or if the fault is unrecoverable, enter one of the following commands:

```
# lpadmin -p printer-name -F continue ↵
# lpadmin -p printer-name -F beginning ↵
# lpadmin -p printer-name -F wait ↵
# lpadmin -p printer-name -F abort ↵
```

These commands direct the LP print service, respectively, to continue at the top of the page, restart from the beginning, wait for you to enter an **enable** command to re-enable the printer (see "Enabling and Disabling Printer" for information on the **enable** command), or abort the job.

If you do not specify how the LP print service is to resume after a printer fault is cleared, it will try to continue at the top of the page where printing stopped, or, failing that, at the beginning of the print request.

If the recovery is **continue**, but the interface program does not stay running so that it can detect when the printer fault has been cleared, printing will be attempted every few minutes until it succeeds. You can force the LP print service to retry immediately by issuing an **enable** command.

User Access Restrictions

You can control which users are allowed to use a particular printer on your system. For instance, if you're designating one printer to handle sensitive information, you don't want all users to be able to use the printer. Another time you might want to do this is when you're authorizing the use of a high quality printer which produces expensive output.

The LP print service will use a list of users allowed or denied access to a printer. The LP print service will reject a user's request to print a file on a printer he or she is not allowed to use. If your users have access to remote printers, or if users on other systems have access to printers on your system, make sure that the allow and deny lists for those printers on your computer match the allow and deny lists on the remote system where the remote printers reside. If these two sets of lists don't match, your users may receive conflicting messages (some accepting jobs, and others refusing jobs) when submitting requests to remote printers.

The method of listing the users allowed or denied access to a printer is similar to the method used to list users allowed or denied access to the **cron** and **at** facilities, and the method described above in "Forms Allowed." Briefly, the rules are as follows:

- An allow list is a list of users allowed to use the printer. A deny list is a list of users denied access to the printer.
- If the allow list is not empty, only the users listed are allowed; the deny list is ignored. If the allow list is empty, users listed in the deny list are not allowed. If both lists are empty, there are no restrictions on who may use the printer.
- Specifying **all** in the allow list allows everybody access to the printer; specifying **all** in the deny list denies access to everybody except the user **lp** and the super-user (**root** or **sysadm**).

You can add names of users to either list using one of the following commands:

```
# lpadmin -p printer-name -u allow:user-list ↵
# lpadmin -p printer-name -u deny:user-list ↵
```

The *user–list* is a comma or space separated list of names of users. If you use spaces to separate the names, enclose the entire list (including the **allow:** or **deny:** but not the **–u**) in quotes. Each item in the *user–list* may take any of the following forms:

user	*user* on any system
all	All users on all systems
system!*user*	*user* on *system* only
!*user*	*user* on local system only
all!*user*	*user* on any system
all!all	All users on all systems
system!**all**	All users on *system*
!all	All users on local system

The first command shown above adds the names to the allow list and removes them from the deny list. The second command adds the names to the deny list and removes them from the allow list.

If you do not use this option, the LP print service will assume that everybody may use the printer.

Inclusion of Banner Page in Output

Most users want to have the output of each print request preceded by a banner page. A banner page shows who requested the printing, the request ID for it, and when the output was printed. It also allows for an optional title that the requester can use to better identify a printout. Finally, the banner page greatly eases the task of separating a sequence of print requests so that each may be given to the correct user.

Sometimes a user needs to avoid printing a banner page. The likely occasions are when the printer has forms mounted that should not be wasted, such as payroll checks or accounts payable checks. Printing a banner page under such circumstances may cause problems.

Enter the following command to allow users to request no banner page with the **lp –o nobanner** command:

```
# lpadmin -p printer-name -o nobanner ↵
```

If you later change your mind, you can reverse this choice by entering the following command:

```
# lpadmin -p printer-name -o banner ↵
```

If you do not allow a user to skip the banner page, the LP print service will reject all attempts to avoid a banner page when printing on the printer. This is the default action.

If you want to disable banner page printing for all print requests on a specific printer, enter the following command:

```
# lpadmin -p printer-name -o banneroff ↵
```

Printer Description

An easy way to give users of the LP print service helpful information about a printer is by adding a description of it. This description can contain any message you like, including the number of the room where the printer is found, the name of the person to call with printer problems, and so forth.

Users can see the message when they use the **lpstat –D –p** *printer–name* command.

To add a description of a printer, enter the following command.

```
# lpadmin -p printer-name -D 'text' ↵
```

The *text* is the message. You'll need to include the quotes if the message contains blanks or other characters that the shell might interpret if the quotes are left out.

Default Printing Attributes

The attributes of a printing job include the page size, print spacing (character pitch and line pitch), and **stty** options for that job. If a user requests a job to be printed on a particular form, the printing attributes defined for that form will be used for that job. When, however, a user submits a print request without requesting a form, the print service uses one of several sets of default attributes.

- If the user has specified attributes to be used, those attributes will take precedence.
- If the user has not specified attributes, but the administrator has executed the **lpadmin –o** command, the default attributes for that command will take precedence.
- If neither of the above, the attributes defined in the **terminfo** database for the printer being used will take precedence.

The LP print service lets you override the defaults for each printer. Doing so can make it easier to submit print requests by allowing you to designate different printers as having different default page sizes or print spacing. A user can then simply route a file to the appropriate printer to get a desired style of output. For example, you can have one printer dedicated to printing wide (132-column) output, another printing normal (80-column by 66-line) output, and yet another printing letter quality (12 characters per inch, 8 lines per inch) output.

You can independently specify four default settings, page width, page length, character pitch, and line pitch. You can scale these to fit your needs: the first two can be given in characters and lines, or inches or centimeters. The last two can be given as characters and lines per inch or per centimeter. In addition, the character pitch can be specified as **pica** for 10 characters per inch (cpi), **elite** for 12 cpi, or **compressed** for the maximum cpi the printer can provide (up to a limit of 30 cpi).

Set the defaults using one or more of the following commands:

```
# lpadmin -p printer-name -o width=scaled-number ↵
# lpadmin -p printer-name -o length=scaled-number ↵
# lpadmin -p printer-name -o cpi=scaled-number ↵
# lpadmin -p printer-name -o lpi=scaled-number ↵
```

Append the letter **i** to the *scaled–number* to indicate inches, or the letter **c** to indicate centimeters. The letter **i** for character pitch (**cpi**) or line pitch (**lpi**) is redundant. You can also give **pica**, **elite**, or **compressed** instead of a number for the character pitch.

If you don't provide defaults when you configure a printer, then the page size and print spacing will be taken from the data for your printer type in the **terminfo** database. (If you do not specify a printer type, the type will be **unknown**, for which there is an entry in the **terminfo** database.) You can find out what the defaults will be by first defining the printer configuration without providing your own defaults, then using the **lpstat** command to display the printer configuration. The command

```
# lpstat -p printer-name -l ↵
```

will report the default page size and print spacing.

Printer Class Membership

This section does not apply if you are making a remote printer accessible to users on your system. It is occasionally convenient to treat a collection of printers as a single class. The benefit is that a user can submit a file for printing by a member of a class, and the LP print service will pick the first printer in the class that it finds free. This allows faster turn-around, as printers are kept as busy as possible.

Classes aren't needed if the only purpose is to allow a user to submit a print request by type of printer. The **lp –T** *content–type* command allows a user to submit a file and specify its type. The first available printer that can handle the type of the file will be used to print it. The LP print service will avoid using a filter, if possible, by choosing a printer that can print the file directly over one that would need it filtered first. See "Providing Filters" for more information about filters.

Classes do have uses, however. One use is to put into a class a series of printers that should be used in a particular order. If you have a high speed printer and a low speed printer, for instance, you probably want the high speed printer to handle as many print requests as possible, with the low speed printer reserved for use when the other is busy. Because the LP print service always checks for an available printer in the order the printers were added to a class, you could add the high speed printer to the class before the low speed printer, and let the LP print service route print requests in the order you wanted.

Until you add a printer to a class, it doesn't belong to one. If you want to do so, use the following command:

```
# lpadmin -p printer-name -c class-name ↵
```

If the class *class–name* doesn't exist yet, it will be created. If you want to remove a printer from a class without deleting the printer, enter the following command:

```
# lpadmin -p printer-name -r class-name ↵
```

The class name may contain a maximum of fourteen alphanumeric characters and underscores. Class names and printer names must be unique. Because they are, a user can specify the destination for a print request without having to know whether it's a class of printers or a single printer.

System Default Destination

You can define the printer or class to be used to print a file when the user has not explicitly asked for a particular destination and has not set the **LPDEST** shell variable. The printer or class must already exist.

Make a printer or class the default destination by entering the following command:

```
# lpadmin -d printer-or-class-name ↵
```

If you later decide that there should be no default destination, enter a null *printer–or–class–name* as in the following command:

```
# lpadmin -d ↵
```

If you don't set a default destination, there will be none. Users will have to explicitly name a printer or class in each print request, (unless they specify the **–T** *content–type* option) or will have to set the LPDEST shell variable with the name of a destination.

Mounting a Form or Print Wheel

See "Providing Forms" later in this section for information about pre-printed forms. Before the LP print service can start printing files that need a preprinted form or print wheel, you must physically mount the form or print wheel on a printer, and notify the LP print service that you have mounted it. (It is not necessary for a form

to be included on the allow list in order to mount it.) If alerting has been set on the form or print wheel, you will be alerted when enough print requests are queued waiting for it to be mounted. (See "Alerting to Mount a Form" below and "Alerting to Mount a Print Wheel" above.)

When you mount a form you may want to see if it is lined up properly. If an alignment pattern has been defined for the form, you can ask that this be repeatedly printed after you've mounted the form, until you have adjusted the printer so that the alignment is correct.

Mounting a form or print wheel involves first loading it onto the printer and then telling the LP print service that it is mounted. Because it is difficult to do this on a printer that's currently printing, and because the LP print service will continue to print files not needing the form on the printer, you will probably have to disable the printer first. Thus, the proper procedure is to follow these three steps:

1. Disable the printer, using the **disable** command.
2. Mount the new form or print wheel as described immediately after this list.
3. Re-enable the printer, using the **enable** command. (The **disable** and **enable** commands are described in "Enabling and Disabling a Printer.")

First, physically load the new form or print wheel into the printer. Then enter the following command to tell the LP print service it has been mounted.

Enter the following command if you are mounting a form:

```
# lpadmin -p printer-name -M -f form-name -a -o filebreak ↵
```

Enter the following command if you are mounting a print wheel:

```
# lpadmin -p printer-name -M -S print-wheel-name ↵
```

If you are mounting a form with an alignment pattern defined for it, you will be asked to press the Enter key before each copy of the alignment pattern is printed. After the pattern is printed, you can adjust the printer and press the Enter key again. If no alignment pattern has been defined, you won't be asked to press the Enter key. You can drop the **–a** and **–o filebreak** options if you don't want to bother with the alignment pattern.

The **–o filebreak** option tells the LP print service to add a form-feed after each copy of the alignment pattern. The actual control sequence used for the form-feed depends on the printer involved and is obtained from the **terminfo** database. If the alignment pattern already includes a form-feed, leave out the **–o filebreak** option.

To unmount a form, use the following command:

```
# lpadmin -p printer-name -M -f none ↵
```

To unmount a print wheel, use the following command:

```
# lpadmin -p printer-name -M -S none ↵
```

Until you've mounted a form on a printer, only print requests that don't require a form will be printed. Likewise, until you've mounted a print wheel on a printer, only print requests that don't require a particular print wheel will be printed. Print requests that do require a particular form or print wheel will be held in a queue until the form or print wheel is mounted.

Removing a Printer or Class

You can remove a printer or class if it has no pending print requests. If there are pending requests, you have to first move them to another printer or class (using the **lpmove** command), or cancel them (using the **cancel** command).

Removing the last remaining printer of a class automatically removes the class as well. The removal of a class, however, does not cause the removal of printers that were members of the class. If the printer or class removed is also the system default destination, the system will no longer have a default destination.

To remove a printer or class, enter the following command:

```
# lpadmin -x printer-or-class-name ↵
```

If all you want to do is to remove a printer from a class without deleting that printer, enter the following command:

```
# lpadmin -p printer-name -r class-name ↵
```

Putting It All Together

It is possible to add a new printer by completing a number of separate steps, shown in the commands described above. You may find it easier, however, to enter one or two commands that combine all the necessary arguments. Below are some examples.

Example 1

Add a new printer called **lp1** (of the type **455**) on printer port **/dev/tty13**. It should use the standard interface program, with the default page size of 90 columns by 71 lines, and linefeeds should *not* be mapped into carriage return/linefeed pairs.

```
# lpadmin -p lp1 -v /dev/tty13 -T 455 -o "width=90 length=71 stty=-onlcr"↵
```

Example 2

Add a new printer called **laser** on printer port **/dev/tty41**. It should use a customized interface program, located in the directory **/usr/doceng/laser_intface**. It can handle three file types—**i10, i300**, and **impress**—and it may be used only by the users **doceng** and **docpub**. (The following command line is broken over multiple lines for readability.)

```
# lpadmin -p laser -v /dev/tty41 -i /usr/doceng/laser_intface \
      -I "i10,i300,impress" -u "allow:doceng,docpub"  ↵
```

Example 3

When you added the **lp1** printer in the first example, you did not set the alerting. Do this now: have the LP print service alert you—by writing to the terminal on which you are logged in—every 10 minutes after a fault until you fix the problem.

```
# lpadmin -p lp1 -A write -W 10 ↵
```

Examining a Printer Configuration

Once you've defined a printer configuration, you probably want to review it to see if it is correct. If after examining the configuration you find you've made a mistake, just reenter the command that applies to the part that's wrong.

Use the **lpstat** command to examine both the configuration and the current status of a printer. The short form of this command gives just the status; you can use it to see if the printer exists and if it is busy, idle, or disabled. The long form of the command gives a complete configuration listing.

Enter one of the following commands to examine a printer.

```
# lpstat -p printer-name ↵
# lpstat -p printer-name -l ↵
```

(The second command is the long form.) With either command you should see one of the following lines of output.

```
printer printer-name now printing  request-id. enabled since date.

printer printer-name is idle. enabled since date.

printer printer-name disabled since date.
        reason

printer printer-name waiting for auto-retry.
        reason
```

The message `waiting for auto-retry` shows that the LP print service failed in trying to use the printer (because of the *reason* shown), and that it will try again later.

With the long form of the command, you should also see the following output:

```
Form mounted: form-name
Content types: content-type-list
Printer type: printer-type
Description: comment
Connection: connection-info
Interface: pathname
On fault: alert-method
After fault: fault-recovery
Users allowed:
     user-list
```

```
Forms allowed:
     form-list
Banner required
Character sets:
     character-set-list
Default pitch: integer CPI, integer LPI
Default page size: scaled-decimal wide, scaled-decimal long
Default port settings: stty-option-list
```

Making Printers Available

There are two steps in making a printer ready for use after you've defined the printer configuration. First, you must instruct the LP print service to accept print requests for the new printer. To do this, run the **accept** command. Second, you must activate or enable the new printer. To do this, run the **enable** command. These tasks are separate steps because you may have occasion to want to do one but not the other.

Accepting Print Requests for a New Printer

Telling the LP print service to accept print requests for the new printer is done with the **accept** command. You will read more about this command in "Managing the Printing Load," later in the section. For now, all you need to know is that you should enter the following command to let this printer be used.

```
# accept printer-or-class-name ↵
```

As you can see, this command is needed to let the LP print service start accepting print requests for a class, too. To prevent the print service from accepting any more requests, execute the following command.

```
# reject printer-or-class-name ↵
```

Enabling and Disabling a Printer

Because you may want to make sure, before printing begins, that the correct form is loaded in your printer, the correct print wheel or font cartridge is in place, and the printer is on-line, the LP print service will wait for an explicit signal from you before it starts printing files. Once you have verified that all the necessary components are in place, you can request the beginning of printing by issuing the **enable** command for a particular printer, as follows:

```
# enable printer-name ↵
```

If you want to enable several printers simultaneously, list the printers (separating the names with spaces) on the same line as the **enable** command. Don't enclose the list in quotes.

Disabling a printer stops further print requests from being printed. (It does not, however, stop the LP print service from accepting new print requests for the

printer.) From time to time you may want to disable a printer. For example, you may want to interrupt a print request, or you may want to change a form or print wheel, in which case you should disable the printer first. Normally, disabling a printer also stops the request that's currently being printed, placing it back in the queue so it can be printed later. You can, however, have the LP print service wait until the current request finishes, or even cancel the request outright.

To disable a printer, enter one of the following commands:

```
# disable -r "reason"  printer-name ↵
# disable -W -r "reason"  printer-name ↵
# disable -c -r "reason"  printer-name ↵
```

The first command disables the printer, stopping the currently printing request and saving it for printing later. The other commands also disable the printer, but the second one makes the LP print service wait for the current request to finish, while the third cancels the current request. The **–c** and **–W** options are not valid when the **disable** command is run to stop a remote printer because, when run for a remote printer, **disable** stops the transferring (rather than the actual printing) of print requests. The *reason* is stored and displayed whenever anyone checks the status of the printer. You can omit it (and the **–r** option) if you don't want to specify a reason.

Several printers can be disabled at once by listing their names in the same line as the **disable** command. You can only enable or disable local printers; the loading of forms, print wheels, and cartridges in a remote printer and the enabling of that printer are the responsibility of the administrator of the remote system. You can, however, enable or disable the transfer of print requests to the remote system on which a printer is located. Only individual printers can be enabled and disabled; classes cannot.

Allowing Users to Enable and Disable a Printer

You may want to make the **enable** and **disable** commands available for use by other users. This availability is useful, for instance, if you have a small organization where anyone who spots a problem with the printer should be able to disable it and fix the problem. This is *not* a good idea if you want to keep others from interfering with the proper operation of the print services.

If you want to allow others access to the **enable** and **disable** commands, use a standard DG/UX system feature called the *setuid bit*. By assigning ownership of these commands to the user **lp** (this should have been done automatically when you installed the software), and by setting the setuid bit, you can make sure that anyone will be allowed to use the **enable** and **disable** commands. Clearing the bit removes this privilege.

To allow everybody to run **enable** and **disable**, enter the following two commands:

```
# chown lp /usr/bin/enable /usr/bin/disable ↵
# chmod u+s /usr/bin/enable /usr/bin/disable ↵
```

The first command makes the user **lp** the owner of the commands; this step should be redundant, but it is safer to run the command than to skip it. The second command turns on the setuid bit.

To prevent others from running **enable** and **disable**, enter the following command:

```
# chmod u-s /usr/bin/enable /usr/bin/disable ↵
```

Troubleshooting

Here are a few suggestions of what to do if you are having difficulty getting a printer to work.

No Output (Nothing is Printed)

The printer is sitting idle; nothing happens. First, check the documentation that came with the printer to see if there is a self-test feature you can invoke; make sure the printer is working before continuing.

There are three possible explanations when you don't receive any output.

Is the Printer Connected to the Computer?

The type of connection between a computer and a printer may vary. See your printer and computer hardware documentation.

Is the Printer Enabled?

The printer must be enabled in two ways: first, the printer must be turned on and ready to receive data from the computer. Second, the LP print service must be ready to use the printer. If you receive error messages when setting up your printer, follow the "fixes" suggested in the messages. When the printer is set up, issue the commands

```
# accept printer-name ↵
# enable printer-name ↵
```

where *printer-name* is the name you assigned to the printer for the LP print service. Now submit a sample file for printing:

```
# lp -d printer-name file-name ↵
```

Is the Baud Rate Correct?

If the baud rate (the rate at which data is transmitted) is not the same for both the computer and the printer, sometimes nothing will print (see below).

Illegible Output

The printer tries printing, but the output is not what you expected; it certainly isn't readable. There are four possible explanations for this situation:

Is the Baud Rate Correct?

Usually, when the baud rate of the computer doesn't match that of the printer, you'll get some output but it will not look at all like what you submitted for printing. Random characters will appear, with an unusual mixture of special characters and unlikely spacing.

Read the documentation that came with the printer to find out what its baud rate is. It should probably be set at 9600 baud for optimum performance, but

that doesn't matter for now. If it isn't set to 9600 baud, you can have the LP print service use the correct baud rate (by default it uses 9600). If the printer is connected via a parallel port, the baud rate is irrelevant.

To set a different baud rate for the LP print service, enter the following command:

```
# lpadmin -p printer-name -o stty=baud-rate ↵
```

Now submit a sample file for printing (explained earlier in this section).

Is the Parity Setting Correct?

Some printers use a "parity bit" to ensure that the data received for printing has not been garbled in transmission. The parity bit can be encoded in several ways; the computer and the printer must agree on which one to use. If they do not agree, some characters either will not be printed or will be replaced by other characters. Generally, though, the output will look approximately correct, with the spacing of "words" typical for your document and many letters in their correct place.

Check the documentation for the printer to see what the printer expects. The LP print service will not set the parity bit by default. You can change this, however, by entering one of the following commands:

```
# lpadmin -p printer-name -o stty=oddp ↵
# lpadmin -p printer-name -o stty=evenp ↵
# lpadmin -p printer-name -o stty=-parity ↵
```

The first command sets odd parity generation, the second sets even parity. The last command sets the default, no parity.

If you are also setting a baud rate other than 9600, you may combine the baud rate setting with the parity settings, as in the sample command below.

```
# lpadmin -p printer-name -o "stty='evenp 1200'" ↵
```

Are Tabs Set Correctly?

If the printer doesn't expect to receive tab characters, the output may contain the complete content of the file, but the text may appear in a chaotic looking format, jammed up against the right margin (see below).

Is the Printer Type Correct?

See below under "Legible Printing but Wrong Spacing."

Are the Default Port Settings (stty) Correct?

If you are printing binary data, the default port settings should include **–opost** which prevents any output processing by the streams device driver. Alternately, you can include **opost** in your **lp** command line: **lp** *file* **–o –opost.**

Legible Printing, but Wrong Spacing

The output contains all of the expected text and may be readable, but the text appears in an undesirable format: double spaced, with no left margin, run together,

or zigzagging down the page. These problems can be fixed by adjusting the printer settings (if possible) or by having the LP print service use settings that match those of the printer. The rest of this section provides details about solving each of these types of problems.

Double Spaced

Either the printer's tab settings are wrong or the printer is adding a linefeed after each carriage return. (The LP print service has a carriage return added to each linefeed, so the combination causes two linefeeds.) You can have the LP print service not send tabs or not add a carriage return by using the **stty –tabs** option or the **–onlcr** option, respectively.)

```
# lpadmin -p printer-name -o stty=-tabs ↵
# lpadmin -p printer-name -o stty=-onlcr ↵
```

No Left Margin/Runs Together/Jammed Up

The printer's tab settings aren't correct; they should be set every 8 spaces. You can have the LP print service not send tabs by using the **–tabs** option.

```
# lpadmin -p printer-name -o stty=-tabs ↵
```

Zigzags Down the Page

The **stty onlcr** option is not set. This is set by default, but you may have cleared it accidentally.

```
# lpadmin -p printer-name -o stty=onlcr ↵
```

A Combination of Problems

If you need to use several of these options to take care of multiple problems, you can combine them in one list, as shown in the sample command below. Include any baud rate or parity settings, too.

```
# lpadmin -p printer-name -o "stty='-onlcr -tabs 2400'" ↵
```

Wrong Characters Printed

If the wrong printer type was selected when you set up the printer with the LP print service, the wrong "control characters" can be sent to the printer. The results are unpredictable and may cause output to disappear or to be illegible, making it look like a problem described above. Another result may be that the wrong control characters cause the printer to set the wrong character set or font.

If you don't know which printer type to specify, try the following to examine the available printer types. First, if you think the printer type has a certain name, try the following command.

```
# tput -T printer-type longname ↵
```

The output of this command will appear on your terminal: a short description of the printer identified by the *printer–type*. Try the names you think might be right until you find one that identifies your printer.

If you don't know what names to try, you can examine the **terminfo** directory to see what names are available. Warning: there are probably many names in that directory. Enter the following command to examine the directory.

```
# ls -CR /usr/share/lib/terminfo/* | more ↵
```

Pick names from the list that match one word or number identifying your printer. Try each of the names in the other command above.

When you have the name of a printer type you think is correct, set it in the LP print service by entering the following command:

```
# lpadmin -p printer-name -T printer-type ↵
```

Dial Out Failures

The LP print service uses UUCP to handle dial-out printers. If a dialing failure occurs and you are receiving printer fault alerts, the LP print service reports the same error reported by UUCP for similar problems. (If you haven't arranged to receive fault alerts, they are mailed, by default, to the user **lp**.) See Appendix A for a list of UUCP error messages and explanations.

Idle Printers

There are several reasons why you may find a printer idle and enabled but with print requests still queued for it:

- The print requests need to be filtered. Slow filters run one at a time to avoid overloading the system. Until a print request has been filtered (if it needs slow filtering), it will not print. Use the following command to see if the first waiting request is being filtered.

  ```
  # lpstat -o -l ↵
  ```

- The printer has a fault. After a fault has been detected, printing resumes automatically, but not immediately. The LP print service waits about five minutes before trying again, and continues trying until a request is printed successfully. You can force a retry immediately by enabling the printer as follows:

  ```
  # enable printer-name ↵
  ```

- A dial-out printer is busy or doesn't answer, or all dial-out ports are busy. As with automatic continuation after a fault, the LP print service waits five minutes before trying to reach a dial-out printer again. If the dial-out printer can't be reached for an hour or two (depending on the reason), the LP print service finally alerts you to a possible problem. You can force a retry immediately by enabling the printer as follows:

  ```
  # enable printer-name ↵
  ```

Networking Problems

You may encounter several types of problems while trying to get files printed over a network: (1) requests being sent to remote printers may back up in the local queue; (2) requests sent to remote printers may be backed up in the remote queue; or (3) a user may receive contradictory messages about whether a remote printer has accepted a print request. The rest of this section describes each of these situations and suggests how to resolve them.

Jobs Backing Up in the Local Queue

There are a lot of jobs backing up in the local queue for a remote printer. There are the following possible explanations:

- The remote system is down or the network between the local and remote systems is down. To resolve this problem, run the **reject** command for all the remote printers on your system, as follows:

  ```
  # reject printer-name ↵
  ```

  This will stop new requests for those printers from being added to the queue. Once the system comes up again, and jobs start being taken from your queue, type **accept** *printer* to allow new jobs to be queued.

- The remote printer is disabled on the local system. The underlying DG/UX Release 5.4 network software was not set up properly. For details, see **lpsystem**(1M).

- The **tcp** listen service does not exist for **lp**. Use the **sysadm** operation Device –> Port –> Port Service –> List to check for the listen service. If the service does not exist, make sure that **/etc/saf/_pmtab** contains the listen service entries for **lp** and **lpr**. Examples of these entries exist in **/etc/saf/_pmtab.proto**.

- The **tcp** listen service is disabled. Enter the **sacadm –l** command on the server to check the status of the port monitor. If the **tcp** status is not "ENABLED," enter the following commands to restart the port monitor:

  ```
  # sacadm -k -p tcp ↵
  # sacadm -s -p tcp ↵
  ```

  Wait for **sacadm –l** to show the monitor is enabled before trying to print.

Jobs Backing Up in the Remote Queue

The remote printer has been disabled.

Conflicting Acceptance/Rejection Messages

A user enters a print request and is notified that the system has accepted it. The job is sent to a remote system and the user receives mail that the job has been rejected.

This may be happening for one of two reasons:

- the local host may be accepting requests while the remote host is rejecting them, or
- the printer's definition on the local host may not match its definition on the remote host.

The definitions of job components (such as filters, character sets, print wheels, and forms) must be identical on both the local and the remote systems if local users are to be able to access remote printers.

Providing Forms

A form is a sheet of paper, on which text or graphical displays have already been printed, that can be loaded into a local printer (that is, a printer on your system) for use in place of plain stock. Common examples of forms include company letterhead, special paper stock, invoices, blank checks, vouchers, receipts, and labels.

Typically, several copies of a blank form are loaded into a printer, either as a tray of single sheets or as a box of fan-folded paper. An application is used to generate data that will be printed on the form, thereby filling it out.

The LP print service helps you manage the use of preprinted forms, but does not provide your application any help in filling out a form; this is solely your application's responsibility. The LP print service, however, will keep track of which print requests need special forms mounted and which forms are currently mounted. It can alert you to the need to mount a new form.

This section tells you how you can manage the use of preprinted forms with the LP print service. You will see how you can

- define a new form
- change the print service's description of an existing form
- remove the print service's description of a form
- examine the print service's description of a form
- restrict user access to a form
- arrange alerting to the need to mount a form
- inform the print service that a form has been mounted

Defining a Form

When you want to provide a new form, the first thing you have to do is define its characteristics. To do so, enter information about each of the nine required characteristics (page length, page width, and so on) as input to the **lpforms** command (see below for details). The LP print service will use this information for two purposes: to initialize the printer so that printing is done properly on the form, and to send you reminders about how to handle that form. Before running the **lpforms** command, gather the following information about your new form:

Page length The length of the form, or of each page in a multi-page form. This can be expressed as the number of lines, or the size in inches or centimeters.

Page width The width of the form, expressed in characters, inches, or centimeters.

Number of pages
The number of pages in a multi-page form. The LP print service uses this number with a filter (if available) to restrict the alignment pattern to a length of one form. (See the description of alignment patterns below.) If no filter is available, the LP print service does not truncate the output.

Line pitch A measurement that shows how closely together separate lines appear on the form. It can be expressed in either lines per inch or lines per centimeter.

Character pitch A measurement that shows how closely together separate characters appear on the form. It can be expressed in either characters per inch or characters per centimeter.

Character set choice
The character set, print wheel, or font cartridge that should be used when this form is used. A user may choose a different character set for his or her own print request when using this form, or you can insist that only one character set be used.

Ribbon color If the form should always be printed using a certain color ribbon, then the LP print service can remind you which color to use when you mount the form.

Comment Any comment you wish to make about the form. This comment is available for users to see so they can understand what the form is, when it should be used, and so on.

Alignment pattern
A sample file that the LP print service uses to fill one blank form. When mounting the form, you can print this pattern on the form to align it properly. You can also define a content type for this pattern so that the printer service knows how to print it.

The LP print service does not try to mask sensitive information in an alignment pattern. If you do not want sensitive information printed on sample forms (very likely the case when you align checks, for instance) then you should mask the appropriate data. The LP print service keeps the alignment pattern stored in a safe place, where only you (that is, the user **lp** and the super-user **root** or **sysadm**) can read it.

When you've gathered this information about the form, enter it as input to the **lpforms** command. You may want to record this information first in a separate file so you can edit it before entering it with **lpforms**. You can then use the file as input instead of typing each piece of information separately after a prompt. Whichever method you use, enter the information in the following format:

```
Page length: scaled-number
Page width: scaled-number
```

```
Number of pages: integer
Line pitch: scaled-number
Character pitch: scaled-number
Character set choice: character-set-name[,mandatory]
Ribbon color: ribbon-color
Comment:
comment
Alignment pattern: [content-type]
alignment-pattern
```

Although these attributes are described in detail on the previous page, a few points should be emphasized here. First, the phrase **mandatory** is optional and, if present, means that the user cannot override the character set choice in the form.

Second, **content–type** can be given optionally, with an alignment pattern. If this attribute is given, the print service uses the alignment pattern specified to determine, as necessary, how to filter and print the file.

With two exceptions, the information in the above list may appear in any order. The exceptions are the alignment pattern (which must always appear last) and *comment* (which must always follow the line with the **Comment:** prompt). If the *comment* contains a line beginning with a key phrase (such as **Page length**, **Page width**, and so on), precede that line with a **>** character so the key phrase is hidden. Be aware, though, that any initial **>** will be stripped from the comment when it is displayed.

Not all of the information has to be given. The defaults for the form's various attributes appear below:

Page length	66 lines
Page width	80 columns
Number of pages	1
Line pitch	6 per inch
Character pitch	10 per inch
Character set choice	any
Ribbon color	any
Comment	(no default)
Alignment pattern	(no default)

To define the form, use one of the following commands

```
# lpforms -f form-name -F file-name ↵
# lpforms -f form-name - ↵
```

where *file–name* is the full path for the file.

The first command gets the form definition from a file; the second command gets the form definition from you, through the standard input. A *form–name* can be anything

you choose, as long as it contains a maximum of fourteen alphanumeric characters and underscores.

If you need to change a form, just reenter one of the above commands. You need only provide information for items that must be changed; items for which you don't specify new information will stay the same.

Removing a Form

The LP print service imposes no fixed limit on the number of forms you may define. It is a good idea, however, to remove forms that are no longer appropriate. If you don't, users will see a long list of obsolete forms when choosing a form, and may be confused. In addition, because the LP print service must occasionally look through all the forms listed before performing certain tasks, the failure to remove obsolete forms may require extra, unnecessary processing by the print service.

To remove a form, enter the following command:

```
# lpforms -f form-name -x ↵
```

Restricting User Access

If your system has a form that you don't want to make available to everyone, you can limit its availability to selected users. For example, you may want to limit access to checks to the people in the payroll department or accounts payable department.

The LP print service restricts the availability of a form by using the list of users allowed or denied access to that form. If a user is not allowed to use a particular form, the LP print service will reject his or her request to print a file with it.

The method used to allow or deny users access to a form is similar to the method used to allow or deny users access to the **cron** and **at** facilities. See the manual pages for **crontab**(1) and **cron**(1M). Briefly, the rules are as follows:

- An allow list is a list of users who are allowed to use the form. A deny list is a list of users who are not allowed to use the form.
- If the allow list is not empty, only the users listed are allowed; the deny list is ignored. If the allow list is empty, the users listed in the deny list are not allowed to to use the form. If both lists are empty, there are no restrictions on who may use the form.
- Specifying **all** in the allow list allows everybody to use the form; specifying **all** in the deny list allows no one except the user **lp** and the super-user (**root** or **sysadm**) to use the form.

If users on your system are to be able to access forms on a remote printer, it's necessary for all the users included on the allow list for the local system to be included on the allow list for the remote system, as well.

If, on the other hand, a local user is to be denied permission to use forms on a remote printer, it's not necessary for the deny lists on both the local and remote

print services to include that user. By being included in only one of these deny lists, a user can be denied access to remote forms. As a courtesy to your users, however, it's a good idea to make sure that any local users who are included in a deny list on a remote system are included in the corresponding deny list on your local system. By doing this you can make sure that whenever a user on your system requests a form that he or she is not authorized to use, he or she is immediately informed that permission to use the form is being denied. If the local print service does not "know" that a user is denied permission to use a particular remote form, there will be a delay before the user receives a "permission denied" message from the remote system.

You can add names of users to either list using one of the following commands:

```
# lpforms -f form-name -u allow:user-list ↵
# lpforms -f form-name -u deny:user-list ↵
```

The *user–list* is a comma or space separated list of names of users. If you use spaces to separate the names, enclose the entire list (including the **allow:** or **deny:** but not the **–u**) in quotes. Each item in the list can include a system name, as shown under "User Access Restrictions" earlier in this section. The first command adds the names to the allow list and removes them from the deny list. The second command adds the names to the deny list and removes them from the allow list. Specifying **allow:all** will allow everybody; specifying **deny:all** will deny everybody.

If you do not add usernames to the allow or deny lists, the LP print service will assume that everybody may use the form.

Alerting to Mount a Form

If you define more forms than printers, you will obviously not be able to print files on all the forms simultaneously. This means that some print requests may be held in a queue until you mount the forms they need. How will you know when to mount a particular form? One method would be to periodically monitor the number of print requests pending for that form. The LP print service, however, provides an easier way: You can ask to be alerted when the number of requests waiting for a form has exceeded a specified threshold.

You can choose one of several ways to receive an alert.

- You can receive an alert via electronic mail. See the manual page for the **mailx**(1) command.
- You can receive an alert written to any terminal on which you are logged in. See the manual page for the **write** command.
- You can receive an alert through a program of your choice.
- You can receive no alerts.

If you elect to receive no alerts, you are responsible for checking to see if any print requests haven't printed because the proper form isn't mounted. In addition to the

method of alerting, you can also set the number of requests that must be queued before you are alerted, and you can arrange for repeated alerts every few minutes until the form is mounted. You can choose the rate of repeated alerts, or choose to receive only one alert for each form.

To arrange for alerting to the need to mount a form, enter one of the following commands:

```
# lpforms -f form-name -A mail -Q requests -W minutes ↵
# lpforms -f form-name -A write -Q requests -W minutes ↵
# lpforms -f form-name -A 'command' -Q requests -W minutes ↵
```

The first two commands direct the LP print service to send you a mail message or write the message directly to your terminal, respectively, for each alert. The third command directs the LP print service to run the *command* for each alert. The shell environment in effect when you enter the third command is saved and restored for the execution of *command*; this includes the environment variables, user and group IDs, and the current directory.

In each command line, the argument *requests* is the number of requests that need to be waiting for the form to trigger the alert, and the argument *minutes* is the number of minutes between repeated alerts. If you want mail sent or a message written to another user when a printer alert occurs, use the third command with the option **–A'mail** *username*' or **–A'write** *username*' .

If you want the print service to issue no alert when the form needs to be mounted, execute the following command:

```
# lpforms -f form-name -A none ↵
```

When you start receiving repeated alerts, you can direct the LP print service to stop sending you alerts (for the current case only) by issuing the following command:

```
# lpforms -f form-name -A quiet ↵
```

Once the form has been mounted and unmounted again, alerts will resume if too many requests are waiting. Alerts will also start again if the number of requests waiting falls below the **–Q** threshold and then rises up to the **–Q** threshold again. This happens when waiting requests are canceled, and when the type of alerting is changed.

If *form–name* is **all** in any of the commands above, the alerting condition applies to all forms for which an alert has not already been defined.

If you don't define an alert method for a form, you will not receive an alert to mount it. If you define a method without the **–W** option, you will be alerted once for each occasion.

Mounting a Form

See "Mounting a Form or Print Wheel" under "Configuring Your Printers" in this section.

Examining a Form

Once you've defined a form to the LP print service, you can examine it with one of two commands, depending on the type of information you want to check. The **lpforms** command displays the attributes of the form. (The display produced by **lpforms** can be used as input; you may want to save it in a file for future reference.) The **lpstat** command displays the current status of the form.

Enter one of the following commands to examine a defined form.

```
# lpforms -f form-name -l ↵
# lpforms -f form-name -l > file-name ↵
# lpstat -f form-name ↵
# lpstat -f form-name -l ↵
```

The first two commands present the definition of the form; the second command captures this definition in a file, which can be used later to redefine the form if you inadvertently remove the form from the LP print service. The last two commands present the status of the form, with the second of the two giving a long form of output, similar to the output of **lpforms –l**:

```
Page length: scaled-number
Page width: scaled-number
Number of pages: integer
Line pitch: scaled-number
Character pitch: scaled-number
Character set choice: character-set[,mandatory]
Ribbon color: ribbon-color
Comment:
comment
Alignment pattern: [ content-type]
content
```

To protect potentially sensitive content, the alignment pattern is not shown if the **lpstat** command is used.

Providing Filters

This section explains how you can manage the use of filters with the LP print service. You will see how you can

- define a new filter
- change a filter
- remove a filter
- examine a filter

The "Customizing the Print Service" section at the end of this section describes how you can write a filter. First, let's see what a filter is and how the LP print service can use one.

What is a Filter?

A filter is a program that you can use for any of three purposes:

- To convert a user's file from one data format to another so that it can be printed properly on a given printer
- To handle the special modes of printing that users may request with the **-y** option to the **lp** command (such as two-sided printing, landscape printing, draft or letter quality printing)
- To detect printer faults and notify the LP print service of them, so that the print service can alert you

Not every filter can perform all three tasks. Given the printer-specific nature of these three roles, the LP print service has been designed so that these roles can be implemented separately. This separation allows you or a printer manufacturer (or another source) to provide filters without having to change the LP print service.

A default filter is provided with the LP print service to provide simple printer fault detection; it does not convert files or handle any of the special modes. It may, however, be adequate for your needs.

Let's examine the three tasks performed by filters more closely.

Task 1: Converting Files

For each printer (local or remote) you can specify what file content types it can print. When a user submits a file to print on any printer, and specifies its content type, the print service will find a printer that can handle files of that content type. Because many applications can generate files for various printers, this is often sufficient. However, some applications may not generate files that can be printed on your printers.

By defining and creating a filter that converts such files into a type that your printers can handle, you can begin to support more applications in the LP print service. (The LP print service comes with a few filters for converting various types of files into PostScript.) For each filter you add to the system, you must specify one or more types of input it can accept and the type of output it can produce (usually only one).

When a user specifies (by executing **lp -T**) a file content type that no printer can handle, the print service tries to find a filter that can convert the file into an acceptable type. If the file to be printed is passed through a filter, the print service will then match the output type of that filter with a printer type or the input type of another filter. The LP print service will continue to match output types to input types in this way, thus passing a file through a series of filters, until the file reaches a printer that accepts it.

Below are some examples.

Example 1
The user Chris has run a spreadsheet program and has generated a file containing a copy of a spreadsheet. Chris now wants to print this file using the LP print service.

You have only HP LaserJet printers on your system. Fortunately, the spreadsheet application understands how to generate output for several printers, and Chris knows it's necessary to request output that can be handled by the HP LaserJet. When Chris submits the file for printing, the LP print service queues it for one of the printers; no filter is needed.

Example 2
Marty has created a graphic image that can be displayed on a Tektronix 4014 terminal. Marty now wants to print this image, but all of the printers are PostScript printers. Fortunately, your system provides a filter called **posttek** that converts **Tektronix** type files to **PostScript**. Because you set the printer type to **PostScript**, the LP print service recognizes that it can use the **posttek** filter to convert Marty's output before printing it.

Task 2: Handling Special Modes

Another important role that filters can perform is the handling of special printing modes. Each filter you add to the filter table can be registered to handle special modes and other aspects of printing:

Special modes
Printer type
Character pitch
Line pitch
Page length
Page width
Pages to print
Character set
Form name
Number of copies

A filter is required to handle the special modes and printing of specific pages; the LP print service provides a default handling for all the rest. However, it may be more efficient to have a filter handle some of the rest, or it may be that a filter has to know several of these aspects to fulfill its other roles properly. A filter may need to know, for example, the page size and the print spacing if it is going to break up the pages in a file to fit on printed pages. As another example, some printers can handle multiple copies more efficiently than the LP print service, so a filter that can control the printer can use the information about the number of copies to skip the LP print service's default handling of multiple copies.

We'll see below how you can register special printing modes and other aspects of printing with each filter.

Task 3: Detecting Printer Faults

Just as converting a file and handling special printing modes is a printer-specific role, so is the detecting of printer faults. The LP print service attempts to detect faults in general, and for most printers it can do so properly. The range of faults that the print service can detect by itself, however, is limited. It can check for hang-ups (loss of carrier, the signal that indicates the printer is on-line) and excessive delays

in printing (receipt of an XOFF flow-control character to shut off the data flow, with no matching XON to turn the flow back on). However, the print service can't determine the cause of a fault, so it can't tell you what to look for.

A properly designed filter can provide better fault coverage. Some printers are able to send a message to the host describing the reason for a fault. Others indicate a fault by using signals other than the dropping of a carrier or the shutting off of data flow. A filter can serve you by detecting more faults and providing more information about them than you would otherwise receive.

Another service a filter can provide is to wait for a printer fault to clear and then to resume printing. This service allows for more efficient printing when a fault occurs because the print request that was interrupted does not have to be reprinted in its entirety. Only a real filter, which has knowledge of the control sequences used by a printer, can "know" where a file breaks into pages; thus only such a filter can find the place in the file where printing should resume.

The LP print service has a simple interface that allows a filter to send you fault information and to restart printing if it can. The alerting mechanism (see "Printer Fault Alerting" under "Configuring Your Printers" in this section) is handled by the LP print service; the interface program that manages the filter takes all error messages from the filter and places them in an alert message that can be sent to you. Thus you'll see any fault descriptions generated by the filter. If you've set the printer configuration so that printing should automatically resume after a fault is cleared, the interface program will keep the filter active, so that printing can pick up where it left off.

Will Any Program Make a Good Filter?

It is tempting to use a program such as **troff**, **nroff**, or a similar word-processing program as a filter. However, the **troff** and **nroff** programs have a feature that allows references to be made in a source file to other files, known as *include files*. The LP print service does not recognize include files; it will not queue any that are referenced by a source file when that file is in a queue to be printed. As a result, the **troff** or **nroff** program, unable to access the include files, may fail. Other programs may have similar features that limit their use as filters.

Here are a few guidelines for evaluating a program for use as a filter:

- Only programs capable of reading data from standard input and writing data to standard output may be used as filters.

- Examine the kinds of files users will submit for printing that will require processing by the program. If they stand alone (that is, if they do not reference other files that the program will need), the program is probably okay.

  Check also to see if the program expects any files other than those submitted by a user for printing. If it does, those files must be in the directory of the person using the filter, and they must be readable by all users authorized to use the filter. The latter prerequisite is necessary because filters are run with the user ID and group ID of the user who submitted the print request.

- If referenced files are permitted in the files submitted for printing, or if the program will need files other than those submitted by a user, then the program, unable to access the additional files, is likely to fail. We suggest you don't use the program under consideration as a filter; instead, have users run the program before submitting files for printing.

Referenced files that are always specified by full pathnames *may* be okay, but only if the filter is used for local print requests. When used on requests submitted from a remote system for printing on your system, the filter may still fail if the referenced files exist only on the remote system.

Filters for Your System

The LP print service is delivered with several filters. As you add, change, or delete filters, you may overwrite or remove some of these original filters. If necessary, you can restore the original set of filters (and remove any filters you have added), with the following command:

```
# lpfilter -f all -i ↵
```

Defining a Filter

When adding a new filter, the first thing you must do is to define the characteristics of its use. To do this, issue the **lpfilter** command with arguments that specify the values of the following filter characteristics:

- the name of the filter (that is, a command name)
- the types of input it will accept
- the types of output it will produce
- the types of printers to which it will be able to send jobs
- the names of specific printers to which it will send jobs
- the type of the filter (whether it's a **fast** filter or a **slow** filter)
- options

Each of these characteristics is described below.

Command:
This is the full path of the filter program. If there are any fixed options that the program always needs, include them here.

Input types:
This is the list of file content types that the filter can process. The LP print service doesn't impose a limit on the number of input types that can be accepted by a filter, but most filters can take only one. Several file types may be similar enough so that the filter can deal with them. You can use whatever names you like here, using a maximum of fourteen alphanumeric characters and dashes (not underscores). Because the LP print service uses

these names to match a filter with a file type, you should follow a consistent naming convention. For example, if more than one filter can accept the same input type, use the same name for that input type when you specify it for each filter. These names should be advertised to your users so they know how to identify the type of a file when submitting that file for printing.

Output types:

This is the list of file types that the filter can produce as output. For each input type the filter will produce a single output type, of course; the output type may vary, however, from job to job. The names of the output types are also restricted to fourteen alphanumeric characters and dashes.

These names should either match the types of printers you have on your system, or match the input types handled by other filters. The LP print service groups filters together in a shell pipeline if it finds that several passes by different filters are needed to convert a file. It's unlikely that you will need this level of sophistication, but the LP print service allows it. Try to find a set of filters that takes (as input types) all the different files your users may want printed, and converts those files directly into types your printers can handle.

Printer types:

This is a list of printer types into which the filter can convert files. For most filters this list will be identical to the list of output types, but it can be different.

For example, you may have a printer that is given a single type for purposes of initialization (see "Printer Type" under "Configuring Your Printers" in this section), but which can recognize several different types of files. In essence this printer has an internal filter that converts the various types into one with which it can deal. Thus, a filter may produce one of several output types that match the "file types" that the printer can handle. The filter should be marked as working with that printer type.

As another example, you may have two different models of printers that are listed as accepting the same types of files. However, due to slight differences in manufacture, one printer deviates in the results it produces. You label the printers as being of different printer types, say A and B, where B is the one that deviates. You create a filter that adjusts files to account for the deviation produced by printers of type B. Because this filter is needed only for those printer types, you would list it as working only on type B printers.

For most printers and filters you can leave this part of the filter definition blank.

Printers:

You may have some printers that, although they're of the correct type for a filter, are in other ways not adequate for the output that the filter will produce. For instance, you may want to dedicate one printer for fast turn-around; only files that the printer can handle without filtering will be sent to that printer. Other printers, of identical type, you allow to be used for files that may need extensive filtering before they can be printed. In this

case, you would label the filter as working with only the latter group of printers.

In most cases a filter should be able to work with all printers that accept its output, so you can usually skip this part of the filter definition.

Filter type:

The LP print service recognizes **fast** filters and **slow** filters. Fast filters are labeled **fast** because they incur little overhead in preparing a file for printing, and because they must have access to the printer when they run. A filter that is to detect printer faults has to be a fast filter.

Slow filters are filters that incur a lot of overhead in preparing a file and that don't require access to a printer. The LP print service runs slow filters in the background, without tying up a printer. This allows files that don't need slow filtering to move ahead; printers will not be left idle while a slow filter works on a file if other files can be printed simultaneously.

Slow filters that are invoked by modes (via the **-y** option), must be run on the computer where the print request was issued. The LP print service can't pass values for modes to remote systems. It can, however, match a file content type (specified after the **-T** option of the **lp** command) to a content type on a remote system. Therefore, if you want to activate special modes on a remote system, you must do so by specifying content types that will allow the LP print service to match input types and output types.

Options:

Options specify how different types of information should be transformed into command line arguments to the filter command. This information may include specifications from a user (with the print request), the printer definition, and the specifications implemented by any filters used to process the request.

There are thirteen sources of information, each of which is represented by a *keyword*. Each option is defined in a *template*, a statement in the following format: *keyword pattern=replacement*. This type of statement is interpreted by the **lpfilter** command to mean "When the information referred to by *keyword* has the value matched by *pattern*, take the *replacement* string, expand any regular expressions it contains, and append the result to the command line."

The options specified in a filter definition may include none, all, or any subset of these thirteen keywords. In addition, a single keyword may be defined more than once, if multiple definitions are required for a complete filter definition. (See "Defining Options with Templates" below.)

When you've gathered enough information to define the above characteristics of your filter, you are ready to run the **lpfilter** command, using your data as arguments. Because there are so many arguments, and because some of them may need to be entered more than once (with different values), we recommend you record this information first in a separate file and edit it, if necessary. You can then use the file as input to the **lpfilter** command and avoid typing each piece of information separately.

Whether you store the information in a file or enter it directly on the command line, use the following format:

```
Command: command-pathname [options]
Input types: input-type-list
Output types: output-type-list
Printer types: printer-type-list
Printers: printer-list
Filter type: fast or slow
Options: template-list
```

The information can appear in any order. Not all the information has to be given. When you do not specify values for the items appearing in Table 12–3, the values shown beside them are assigned by default.

Table 12–3 Filter Defaults

Item	Default
Command:	(no default)
Input types:	any
Output types:	any
Printer types:	any
Printers:	any
Filter type:	slow
Options:	(no default)

As you can see, the default values define a very flexible filter, so you probably have to supply at least the input and output types. When you enter a list, you can separate the items in it with blanks or commas, unless it is a *templates–list*; items in a *templates–list* must be separated by commas.

Defining Options with Templates

A template is a statement in a filter definition that defines an option to be passed to the filter command based on the value of one of the characteristics of the filter. A filter definition may include more than one template. Multiple templates may be entered on a single line and separated with commas, or they may be entered on separate lines, preceded by the **Options:** prefix.

The format of a template is as follows:

```
keyword pattern = replacement
```

The *keyword* identifies the type of option being registered for a particular characteristic of the filter.

Let's look at an example of how an option is defined for a particular filter. Suppose you want to have the print service scheduler assign print requests to filters on the basis of the following criteria:

- If the type of output to be produced by the filter is **impress**, then pass the **–I** option to the filter.
- If the type of output to be produced by the filter is **postscript**, then pass the **–P** option to the filter.

To specify these criteria, provide the following templates as options to the lpfilter command.

```
Options: OUTPUT impress=-I, OUTPUT postscript=-P
```

If the **Options:** line becomes too long, put each template on a separate line, as follows:

```
Options: OUTPUT impress=-I
Options: OUTPUT postscript=-P
```

In both templates, the *keyword* is defined as **OUTPUT**. In the first template, the value of *pattern* is **impress** and the value of the *replacement* is **–I**. In the second template, the value of *pattern* is **postscript** and the value of the *replacement* is **–P**.

Template Keywords

Table 12–4 shows the thirteen keywords available for defining options in a filter definition.

Table 12–4 Filter Option Keywords

Characteristic	Keyword	Possible Patterns	Example
Content type (input)	INPUT	content-type	troff
Content type (output)	OUTPUT	content-type	postscript
Printer type	TERM	printer-type	att495
Printer name	PRINTER	printer-name	lp1
Character pitch	CPI	scaled-decimal	10
Line pitch	LPI	scaled-decimal	6
Page length	LENGTH	scaled-decimal	66
Page width	WIDTH	scaled-decimal	80
Pages to print	PAGES	page-list	1–5,13–20
Character set	CHARSET	character–set	finnish
Form name	FORM	form-name	invoice2
Number of copies	COPIES	integer	3
Special modes	MODES	mode	landscape

To find out which values to supply for each type of template (that is, for the *pattern* and *replacement* arguments for each *keyword*), see the source of information listed below.

- The values for the **INPUT** and **OUTPUT** templates come from the file type that needs to be converted by the filter and the output type that has to be produced by the filter, respectively. They'll each be a type registered with the filter.

- The value for the **TERM** template is the printer type.

- The value for the **PRINTER** template is the name of the printer that will be used to print the final output.

- The values for the **CPI**, **LPI**, **LENGTH**, and **WIDTH** templates come from the user's request, the form being used, or the default values for the printer.

- The value for the **PAGES** template is a list of pages that should be printed. Typically, it is a comma separated list of page ranges, each of which consists of a dash separated pair of numbers or a single number (such as **1–5,6,8,10** for pages 1 through 5, 6, 8, and 10). However, whatever value was given in the **–P** option to a print request is passed unchanged.

- The value for the **CHARSET** template is the name of the character set to be used.

- The value for the **FORM** template is the name of the form requested by the **–f** option of the **lp** command.

- The value of the **COPIES** template is the number of copies that should be made of the file. If the filter uses this template, the LP print service will reduce to 1 the number of copies of the filtered file it will print, since this single copy will actually comprise the multiple copies produced by the filter.

- The value of the **MODES** template comes from the **–y** option of the **lp** command (the command used to submit a print request). Because a user can specify several **–y** options, there may be several values for the **MODES** template. The values will be applied in the left-to-right order given by the user.

The *replacement* part of a template shows how the value of a template should be given to the filter program. It is typically a literal option, sometimes with the place-holder * included to show where the value goes. The *pattern* and *replacement* can also use the regular expression syntax of **ed**(1) for more complex conversion of user input options into filter options. All of the regular expression syntax of **ed**(1) is supported, including the **\(** ... **\)** and *****n* constructions, which can be used to extract portions of the *pattern* for copying into the *replacement*, and the **&**, which can be used to copy the entire *pattern* into the *replacement*. If a comma or an equals sign (**=**) is included in a *pattern* or a *replacement*, escape its special meaning by preceding it with a backslash (****). Note that some regular expressions include commas that will have to be escaped this way. A backslash in front of any of these characters is removed when the *pattern* or *replacement* is used.

The following examples show how this works.

Example 1

You provide the following filter definition for a filter called **col**.

```
Input types:  N37, Nlp, simple
Output types:  simple
Command:  /usr/bin/col
Options:  TERM 450 = -b, MODES expand = -x
Options:  INPUT simple = -p -f
```

If you provide more than one definition (that is, more than one line) for any filter characteristic other than **OPTIONS**, only the second definition will be used by the print service.

After you have registered this definition with the print service by entering it as input with the **lpfilter** command, users' print requests will be handled as follows:

If a user enters the command

```
# lp -y expand report.dec10 ↵
```

the filter command will be run with the following arguments:

```
/usr/bin/col -x -p -f
```

If a user enters the command

```
# lp -T N37 -y expand report.dec10 ↵
```

the filter command will be run with the following arguments:

```
/usr/bin/col -x
```

Qualifier: The default printer is not of type **450**.

If a user enters the command

```
# lp -y expand -T 450 report.dec10 ↵
```

the filter command will be run with the following arguments:

```
/usr/bin/col -b -x
```

Example 2

The filter program is called **/usr/lib/lp/postscript/dpost**. It takes one input type, **troff**, produces an output type called **postscript**, and works with any printer of type **postscript**. You've decided that your users need give just the abbreviations **port** and **land** when they ask for the paper orientation to be portrait mode and landscape mode, respectively. Because these options are not intrinsic to the LP print service, users must specify them using the **-y** option to the **lp** command.

The filter definition would look like this:

```
Input types: troff
Output types: postscript
Printer types: postscript
Filter type: slow
Command: /usr/lib/lp/postscript/dpost
Options: LENGTH * = -l *, CHARSET * = -s *
Options: MODES port = -o portrait, MODES land = -o landscape
```

A user submitting a file of type **troff** for printing on a PostScript printer (type **postscript**), with requests for landscape orientation and the **gothic** character set, would enter the following command:

```
# lp -T troff -S gothic -y land -d any ↵
```

Then this filter would be invoked by the LP print service to convert the file as follows:

```
/usr/lib/lp/postscript/dpost -S gothic -o landscape
```

Example 3
You add the following option template to the previous example:

```
Options:  MODES size\=\([0-9]*\)x\([0-9]*\) = -h\1 -w\2
```

This template is used to convert a **MODES** option of the form

-y size=*height***x***width*

into a pair of filter options,

-h*height* -w*width*

So if a user gives the following command

```
# lp -y size=24x80 ↵
```

the **dpost** command would include the following options:

```
-h24 -w80
```

Command to Enter

Once a filter definition is complete, enter one of the following commands to add the filter to the system.

 # lpfilter -f *filter-name* **-F** *file-name* ↵
 # lpfilter -f *filter-name* **-** ↵

The first command gets the filter definition from a file, and the second command gets the filter definition from the standard input. A *filter–name* can be any string you choose, with a maximum of fourteen alphanumeric characters and underscores.

If you need to change a filter, just reenter one of the same commands. You need only provide information for those items that must be changed; items for which you don't specify new information will stay the same.

Removing a Filter

The LP print service imposes no fixed limit on the number of filters you can define. It is a good idea, however, to remove filters no longer applicable, to avoid extra

processing by the LP print service which must examine all filters to find one that works in a given situation.

To remove a filter, enter the following command:

```
# lpfilter -f filter-name -x ↵
```

Examining a Filter

Once you've added a filter definition to the LP print service, you can examine it by running the **lpfilter** command. The output of this command is the filter definition displayed in a format that makes it suitable as input. You may want to save this output in a file that you can use later to redefine the filter if you inadvertently remove the filter from the LP print service.

To examine a defined filter, enter one of the following commands:

```
# lpfilter -f filter-name -l ↵
# lpfilter -f filter-name -l >file-name ↵
```

The first command presents the definition of the filter on your screen; the second command captures this definition in a file for future reference.

Restoring Factory Defaults

The software is shipped from the factory with default values set for all options. If, after changing these defaults, you want to reset them, execute the following command:

```
# lpfilter -f filter-name -i ↵
```

You can restore an individual filter by giving its name in place of *filter–name*, or you can restore everything by giving the name **all**.

A Word of Caution

Adding, changing, or deleting filters can cause print requests still queued to be canceled. This is because the LP print service evaluates all print requests still queued, to see which are affected by the filter change. Requests that are no longer printable, because a filter has been removed or changed, are canceled (with notifications sent to the people who submitted them). There can also be delays in the responses to new or changed print requests when filters are changed, due to the many characteristics that must be evaluated for each print request still queued. These delays can become noticeable if there is a large number of requests that need to be filtered.

Because of this possible impact, you may want to make changes to filters during periods when the LP print service is not being used much.

Managing the Printing Load

Occasionally you may need to stop accepting print requests for a printer or move pending print requests from one printer to another. There are various reasons why you might want to do this, such as the following:

- the printer needs periodic maintenance
- the printer is broken
- the printer has been removed
- you've changed the configuration so that the printer is to be used differently
- too many large print requests are queued for one printer and should be spread around

If you are going to make a big change in the way a printer is to be used, such as stopping its ability to handle a certain form, changing the print wheels available for it, or disallowing some users from using it, print requests that are currently queued for printing on it will have to be moved or canceled. The LP print service will attempt to find alternate printers, but only if the user doesn't care which printer is to be used. Requests for a specific printer won't be automatically moved; if you don't move them first, the LP print service will cancel them.

If you decide to take a printer out of service, to change its configuration, or to lighten its load, you may want to move print requests off it and reject additional requests for it for awhile. To do so, use the **lpmove** and **reject** commands. If you do reject requests for a printer, you can accept requests for it later, by using the **accept** command.

Rejecting Requests for a Printer or Class

To stop accepting any new requests for a printer or class of printers, enter the following command.

```
# reject -r "reason"  printer-or-class-name ↵
```

You can reject requests for several printers or classes in one command by listing their names on the same line, separating the names with spaces. The *"reason"* will be displayed whenever anyone tries to print a file on the printer. You can omit it (and the **-r**) if you don't want to specify a reason.

Although the **reject** command stops any new print requests from being accepted, it will not move or cancel any requests currently queued for the printer. These will continue to be printed as long as the printer is enabled.

Accepting Requests for a Printer or Class

After the condition that led to rejecting requests has been corrected or changed, enter one of the the following commands to start accepting new requests.

```
# accept printer-name ↵
# accept class-name ↵
```

Again, you can accept requests for several printers or classes in one command by listing their names on the same line.

You will always have to use the **accept** command for a new printer or class after you have added it, because the LP print service does not initially accept requests for new printers or classes.

Moving Requests to Another Printer

If you specify **–d any** when you run the **lp** command to queue a job, the print service schedules the job for a particular printer. If another becomes available first, the job is sent to the latter printer. If a job is scheduled for a given printer and you run **lpmove** to get jobs off that printer, that job will be moved off and the destination will change from **any** to the printer you've specified on the **lpmove** command line. Users may not have intended this side effect. If not, run the following command:

```
# lp -i request-ID -d any ↵
```

This command will change the destination for the requested job to the original destination: **any** (that is, any available printer).

If you have to move requests from one printer or class to another, enter one of the following commands

```
# lpmove request-id printer-name ↵
# lpmove printer-name1 printer-name2 ↵
```

You can give more than one request ID before the printer name in the first command.

The first command above moves the listed requests to the printer *printer–name*. The second command tries to move *all* requests currently queued for *printer–name1* to *printer–name2*. If some requests cannot be printed on the new printer, they will be left in the queue for the original printer. When the second command is used, the LP print service also stops accepting requests for *printer–name1* (the same result you would obtain by running the command **reject** *printer–name2*).

Examples

Here are some examples of how you might use these three commands.

Example 1

You've decided it is time to change the ribbon and perform some preventive maintenance on printer **lp1**. First, to prevent the loss of print requests already queued for **lp1**, you move all requests from printer **lp1** to printer **lp2**.

```
# lpmove lp1 lp2 ↵
```

After the requests are moved, make sure the LP print service does not print any more requests on **lp1** by disabling it.

```
# disable lp1 ↵
```

Now you may physically disable the printer and start working on it.

Example 2

You've finished changing the ribbon and doing the other work on **lp1**; now it's time to bring it back into service. Execute the following commands in any order:

```
# accept lp1 ↵
# enable lp1 ↵
```

See "Enabling and Disabling a Printer" under "Making Printers Available" in this section.

Example 3

You notice that someone has queued several large files for printing on the printer **laser1**. Meanwhile **laser2** is idle because no one has queued requests for it. Move the two biggest requests (**laser1–23** and **laser1–46**) to **laser2**, and reject any new requests for **laser1** for the time being.

```
# lpmove  laser1-23  laser1-46  laser2  ↵
# reject  -r"too busy--will reopen later"  laser1  ↵
```

Managing Queue Priorities

The LP print service provides a simple priority mechanism that users can use to adjust the position of a print request in the queue. Each print request can be given a priority level by the user who submits it; this is a number from 0 to 39, with *smaller* numbers indicating *higher* levels of priority. Requests with higher priority (smaller numbers) are placed ahead of requests with lower priority (larger numbers).

Thus, for example, a user who decides that her print request is of low priority can assign it a larger value when she submits the file for printing. Another user who decides that his print request is of high priority can assign it a smaller value when he submits the file for printing.

A priority scheme this simple would not work if there were no controls on how high one can set the priority. You can define the following characteristics in this scheme:

- Each user can be assigned a priority limit. One cannot submit a print request with a priority higher than his or her limit, although one can submit a request with a lower priority.
- A default priority limit can be assigned for the balance of users not assigned a personal limit.
- A default priority can be set. This is the priority given print requests to which the user does not assign a priority.

By setting the characteristics according to your needs, you can prevent lower priority printing tasks (such as regular printing by most staff members) from

interfering with higher priority printing tasks (such as payroll check printing by the accounting staff).

You may find that you want a critical print request to print ahead of any others, perhaps even if it has to preempt the currently printing request. You can have the LP print service give immediate handling to a print request, and you can have it hold another print request. This will allow the first request to be printed and will delay the second print request until you allow it to be resumed.

The **lpusers** command lets you assign both priority limits for users and priority defaults. In addition, you can use the **lp –i***request–id* **–Hhold** and **lp –i***request–id* **–Himmediate** commands to put a request on hold or to move it up for immediate printing, respectively. These commands are discussed in detail below.

Setting Priority Limits

To set a user's priority limit, enter the following command.

```
# lpusers -q priority-level -u username ↵
```

You can set the limit for a group of users by listing their names after the **–u** option. Separate multiple names with a comma or space (enclose the list in quotes if you use a space, though). The argument *priority–level* is a number from 0 to 39. As mentioned before, the lower the number the higher the priority, or, in this case, the priority limit.

If you want to set a priority limit for the remaining users, enter the following command:

```
# lpusers -q priority-level ↵
```

This sets the default limit; the default applies to those users for whom you have not set a personal limit, using the first **lpusers** command.

If you later decide that someone should have a different priority limit, just reenter the first command above with a new limit. Or, if you decide that the default limit is more appropriate for someone who already has a personal limit, enter the following command:

```
# lpusers -u username ↵
```

Again, you can do this for more than one user at a time by including a list of names. Using the **lpusers** command with just the **–u** option removes users' personal priority limits and puts the default limit into effect for those users.

Setting a Default Priority

To set the default priority (the priority level assigned to print requests submitted without a priority), use the following command:

```
# lpusers -d priority-level ↵
```

Don't confuse this default with the default limit. This default is applied when a user doesn't specify a priority level; the default limit is applied if you haven't assigned a limit for a user—it is used to limit the user from requesting too high a priority. If the default priority is greater than the limit for a user, the limit is used instead.

If you do not set a default priority, the LP print service will use a default of 20.

Examining the Priority Limits and Defaults

You can examine all the settings you have assigned for priority limits and defaults by entering the following command.

```
# lpusers -l ↵
```

Moving a Request Around in the Queue

Once a user has submitted a print request, you can move it around in the queue to some degree:

- you can adjust the priority to any level, regardless of the limit for the user (who may adjust it only up to his or her limit)
- both you and the user can put it on hold and allow other requests to be printed ahead of it
- you can put it at the head of the queue for immediate printing

Use the **lp**(1) command to do any of these tasks.

Changing the Priority for a Request

If you want to change the priority of a particular request that is still waiting to be printed, you can assign a new priority level to it. By doing so, you can move it in the queue so that it is ahead of lower priority requests, and behind requests at the same level or of higher priority. The priority limit assigned to the user (or the default priority limit) has no effect because, as the administrator, you can override this limit.

Enter the following command to change the priority of a request.

```
# lp -i request-ID -q new-priority-level ↵
```

You can change only one request at a time with this command.

Putting a Request on Hold

Any request that has not finished printing can be put on hold. This will stop its printing, if it is currently printing, and keep it from printing until you resume it. A user may also put his or her own request on hold and then resume it, but may not resume a print request that has been put on hold by the administrator.

To place a request on hold, enter the following command:

```
# lp -i request-ID -H hold ↵
```

Enter the following command to resume the request:

```
# lp -i request-ID -H resume ↵
```

Once resumed a request will continue to move up the queue and will eventually be printed. If printing had already begun when you put it on hold, it will be the next request printed.

Moving a Request to the Head of the Queue

You can move a print request to the head of the queue where it will be the next one eligible for printing. If you want it to start printing immediately but another request is currently being printed, you may interrupt the first request by putting it on hold, as described above.

Enter the following command to move a print request to the head of the queue:

```
# lp -i request-ID -H immediate ↵
```

Only you, as the administrator, can move a request in this way; regular users cannot use the **-H immediate** option. If you set more than one request for immediate printing, the requests will be printed in the reverse order set; that is, the request moved to the head of the queue most recently will be printed first.

Starting and Stopping the LP Print Service

Under normal operation, you should never have to start or stop the LP print service manually. The system starts it each time the system boots, and stops it each time the system stops. If, however, you need to stop the LP print service without stopping the entire system, do so by following the procedure described below.

Stopping the LP print service will cause all printing to cease within seconds. Any print requests that have not finished printing will be printed in their entirety after the LP print service is restarted. The printer configurations, forms, and filters in effect when the LP print service is stopped will be restored after it is restarted. To start and stop the LP print service manually, you must be logged in as either the user **lp** or the super-user (**root** or **sysadm**).

Manually Stopping the Print Service

To stop the LP print service manually, enter the following command:

```
# lpshut ↵
```

The message

```
Print services stopped.
```

appears, and all printing will cease within a few seconds. If you try to stop the LP print service when it is not running, you will see the message

```
Print services already stopped.
```

For additional cleanup and error recovery, you should use the following command to stop the print service:

```
# sysadm stoplp ↵
```

This executes **admlp** scripts which provide additional cleanup features.

Manually Starting the Print Service

To restart the LP print service manually, enter the following command:

```
# lpsched ↵
```

The message:

```
Print services started.
```

will appear. It may take a minute or two for the printer configurations, forms, and filters to be reestablished, before any saved print requests start printing. If you try to restart the LP print service when it is already running, you will see the message

```
Print services already active.
```

For additional cleanup and error recovery, you should use the following command to start the print service:

```
# sysadm startlp ↵
```

This executes **admlp** scripts which provide additional cleanup features.

The LP print service does not have to be stopped to change printer configurations or to add forms or filters.

Managing the LP Print Service Logs

The LP print service has several logs that accumulate information on various LP services. By default, the LP print service does not remove or truncate these logs; therefore, you need to make sure they do not consume too much disk space. The easiest way to manage these logs is by submitting the **cron** jobs provided for the purpose in the LP print service prototype **crontab** file, **/admin/crontabs/lp.proto**. For more information on submitting **cron** jobs, see Chapter 2. Also see the manual pages for **cron**(1M) and **crontab**(1).

Located in **/var/spool/lp/logs**, the LP print service logs are **lpNet**, **lpsched**, and **requests**. The following sections elaborate.

The lpNet Log

The LP system runs a process called **lpNet**, which handles communication with LP systems running on other systems in your network. The **lpNet** process logs its activities to the file **/var/lp/logs/lpNet**.

As the **lpNet** file grows, you should clean it out occasionally. The easiest way to do this is with the appropriate **cron** job from **/admin/crontabs/lp.proto**. The job for this purpose appears below.

```
17 3 * * 0 /usr/lbin/agefile -c4 /var/lp/logs/lpNet
```

This **cron** job ages the **lpNet** log so that every week starts a new **lpNet** log, and no log remains on the system for more than five weeks. This job runs every Sunday at 3:17 am.

The lpsched Log

The LP print service keeps a log for scheduler-related events. This log is **/var/spool/lp/logs/lpsched**.

As with the **lpNet** log, the best way to control the size of the **lpsched** log is by submitting the appropriate **cron** job from **/admin/crontabs/lp.proto**. The job for this purpose appears below.

```
15 3 * * 0 /usr/lbin/agefile -c4 /var/lp/logs/lpsched
```

This **cron** job ages the **lpsched** log so that every week starts a new **lpsched** log, and no log remains on the system for more than five weeks. This job runs every Sunday at 3:15 am.

The requests Log

The **/var/spool/lp/logs/requests** file accumulates information about each completed request. Initially, the directories **/var/spool/lp/tmp/***system* and **/var/spool/lp/requests/***system* contain files that describe each request that has been submitted to the LP print service. Each request has two files (one in each directory) that contain information about the request. The information is split to put more sensitive information in the **/var/spool/lp/requests/***system* directory where it can be kept secure: the request file in the **/var/spool/lp/tmp/***system* is safe from all except the user who submitted the request, while the file in **/var/spool/lp/requests/***system* is safe from all users, including the submitting user.

These files remain in their directories only as long as the request is in the queue. Once the request is finished, the information in the files is combined and appended to the file **/var/lp/logs/requests**.

Log Structure

The **requests** log has a simple structure that makes it easy to extract data from it using common shell commands. Requests are listed in the order they are printed, and are separated by lines showing their request IDs. Each line below the separator line is marked with a single letter that identifies the kind of information contained in that line. Each letter is separated from the data by a single space. See the following list for details.

= This is the separator line. It contains the request ID, the user and group IDs of the user, the total number of bytes in the original (unfiltered) files, and

the time when the request was queued. These items are separated by commas and are in the order just named. The user ID, group ID, and sizes are preceded by the words **uid**, **gid**, and **size**, respectively.

C The number of copies printed.

D The printer or class destination or the word **any**.

F The name of the file printed. This line is repeated for each file printed; files were printed in the order given.

f The name of the form used.

H One of three types of special handling: **resume**, **hold**, and **immediate**. The only useful value found in this line will be **immediate**.

N The type of alert used when the print request was successfully completed. The type is the letter **M** if the user was notified by mail, or **W** if the user was notified by a message to his or her terminal.

O The **–o** options.

P The priority of the print request.

p The list of pages printed.

r This single letter line is included if the user asked for *raw* processing of the files (the **–r** option of the **lp** command).

S The character set or print wheel used.

s The outcome of the request, shown as a combination of individual bits expressed in hexadecimal form. While several bits are used internally by the print service, the most important bits are listed below:

`0x0004` Slow filtering finished successfully.

`0x0010` Printing finished successfully.

`0x0040` The request was canceled.

`0x0100` The request failed filtering or printing.

T The title placed on the banner page.

t The type of content found in the files.

U The name of the user who submitted the print request.

x The slow filter used for the request.

Y The list of special modes to give to the filters used to print the request.

y The fast filter used for the request.

z The printer used for the request. This will differ from the destination (the **D** line) if the request was queued for any printer or a class of printers, or if the request was moved to another destination by the LP print service administrator.

Cleaning out the requests Log

The LP print service does not remove the **requests** file; therefore, the file will grow indefinitely if you do not remove it or shorten it occasionally. The best way to do this is with the **cron** job provided in the LP print service's prototype **crontab** file, **/admin/crontabs/lp.proto**. The prototype file contains three **cron** jobs. The one for managing the **requests** file appears below.

```
13 3 * * * cd /var/lp/logs; if [ -f requests ]; then \
  /bin/mv requests xyzzy; /bin/cp xyzzy requests; >xyzzy; \
  /usr/lbin/agefile -c2 requests; /bin/mv xyzzy requests; fi
```

This jobs appears as one line in the **crontab** file, but it is split into several lines here for readability.

What this entry does, briefly, is *age* the file, changing the name to **requests–1**, and moving the previous day's copy to **requests–2**. The number **2** in the **–c** option to the **agefile** program keeps the log files from the previous two days, discarding older log files. By changing this number you can change the amount of information saved. On the other hand, if you want the information to be saved more often, or if you want the file to be cleaned out more often than once a day, you can change the time when the **crontab** entry is run by changing the first two numbers. The current values, **13** and **3**, cause cleaning up to be done at 3:13 am each day.

The default **crontab** entry supplied is sufficient to keep the old print request records from accumulating in the spooling file system. You may want to condense information in the request log to produce a report on the use of the LP print service, or to aid in generating accounting information. You can produce a different script that examines the file and extracts information just before the clean up procedure.

PostScript Printers

PostScript is a general purpose programming language, like C or Pascal. In addition to providing the usual features of a language, however, PostScript allows a programmer to specify the appearance of both text and graphics on a page.

A PostScript printer is a printer equipped with a computer that runs an interpreter for processing PostScript language files. When a PostScript printer receives a file, it runs that file through the interpreter and then prints it. Unless special provisions have been made by the manufacturer, files submitted to a PostScript printer must be written in the PostScript language.

Why would you want to use a PostScript printer? PostScript provides excellent facilities for managing text and graphics and combining them. Graphics operators facilitate the construction of geometric figures which can then be positioned and scaled with any orientation. The text capabilities allow the user to specify a number of different fonts that can be placed on a page in any position, size, or orientation. Because text is treated as graphics, text and graphics are readily combined. Moreover, the language is resolution and device independent, so that draft copies can be proofed on a low-resolution device and the final version printed in higher resolution on a different device.

Applications that support PostScript, including word-processing and publishing software, will create documents in the PostScript language without intervention by the user. Thus, it is not necessary to know the details of the language to take advantage of its features. However, standard files that many applications produce cannot be printed on a PostScript printer because they are not described in the language. The LP print service provides optional filters to convert many of these files to PostScript so that users may take advantage of PostScript and continue to use their standard applications such as **troff**.

How to Use a PostScript Printer

When the PostScript printers and filters have been installed, the LP subsystem manages PostScript files like any others. If **psfile** is a file containing a PostScript document and **psprinter** has been defined as a PostScript printer, the command

```
# lp -d psprinter -Tpostscript psfile
```

will schedule the print request and manage the transmission of the request to the PostScript printer.

Support of Non-PostScript Print Requests

A PostScript printer may not be able to interpret every kind of file that an application sends to it. The following list names a few of the formats that your printer may not be able to accommodate.

troff Print a **troff** text file.

simple Print an ASCII text file.

dmd Print the contents of a bit-mapped display (for example, an AT&T 630).

tek4014 Print files formatted for a Tektronix 4014 device.

daisy Print files intended for a daisy-wheel printer (for example, a Diablo 630).

plot Print plot-formatted files

Optional filters are provided with the LP print service to translate print requests with these formats to the PostScript language. For example, to convert a file containing ASCII or **troff** code to PostScript code, the filter takes that text and writes a program around it, specifying printing parameters such as fonts and the layout of the text on a page.

Once the PostScript filters are installed, they will be invoked automatically by the LP print service when a user specifies a content-type for a print request with the **-T** option. For example,

```
# lp -d psprinter -T simple report2
```

will automatically convert the ASCII file **report2** (a file with an ASCII or **simple** format) to PostScript when the destination printer *psprinter* has been defined to the system as a PostScript printer. Note that `-T simple` is the default content type and is used here for example purposes only.

Additional PostScript Capabilities Provided by Filters

The filters previously described also take advantage of PostScript capabilities to provide additional printing flexibility. Most of these features may be accessed through the mode option (invoked by the **-y** option) to the **lp** command. These filters allow you to use several unusual options for your print jobs. The following list describes these options and shows the option you should include on the **lp** command line for each one. Most of these options are provided by the **postprint** filter. See the **postprint**(1) manual page for more information.

-y reverse
Reverse the order in which pages are printed.

-y landscape
Change the orientation of a physical page from portrait to landscape.

-y x=*xnumber***,y=***ynumber*
Change the default position of a logical page on a physical page by moving the origin.

-y group=*number*
Group multiple logical pages on a single physical page.

-P *number*
Select, by page numbers, a subset of a document to be printed

If these filters are to be used with an application that creates PostScript output, make sure that the application conforms to the PostScript file structuring comments. In particular, the beginning of each PostScript page must be marked by the comment

```
%%Page:label ordinal
```

where *ordinal* is a positive integer that specifies the position of the page in the sequence of pages in the document.

For example, say you have a file called **report2** that has a content type **simple** (meaning that the content of this file is in ASCII format). You want to print six pages of this file (pages 4-9) in landscape mode (that is, sideways on the page) with two logical pages on each physical page. Because one of the printers on your system (**psprinter**) is a PostScript printer, you can do this, by entering the following command:

```
# lp -d psprinter -T simple -P 4-9 -y landscape,group=2 report2 ↵
```

In addition, the LP print service offers a special filter that can print a gray-scale representation of a matrix. (A gray-scale matrix is a matrix in which each cell is colored one of seven shades of gray to indicate the value of the cell. Darker shades correspond to larger values.) To print a gray-scale representation, specify **matrix** as the content type of your source file by giving the **-T matrix** option.

The Administrator's Duties

Support of PostScript printers is similar to support of other printers, in that the printers must be defined to the system with the **lpadmin** command and the

appropriate software must be installed to manage them. PostScript printers may require some additional effort in supporting fonts.

NOTE: A print job submitted to a serial PostScript printer on an AViiON 500 series system may be truncated. This problem may occur with Data General model 6646T, 6779T, 6772, and 6773 printers that are connected by cable 15340E, through a VDA cluster box (VDC/8p/16) or a VAC 16. This will be fixed in a future release. Meanwhile, to work around the problem, add **clocal** to the "Stty options:" query with the **sysadm** `Device -> Printer -> Device -> Modify` operation.

Installing and Maintaining PostScript Printers

PostScript printers, like other printers, are installed with the **lpadmin** command. The content type of a PostScript printer must be consistent with the content type used in PostScript filters. Consequently, you should install serial PostScript printers with a content type of **PS**, causing LP to use the **postio** filter to communicate with the printer. When installing a parallel printer, however, specify a content type of **postscript** instead, causing LP not to use the **postio** filter. The **postio** filter requires a bidirectional communication connection, which can only be supplied by serial lines.

The content type accepted by the printer is specified with the **–I** option in the **lpadmin** command. In addition, you must tell the LP print service which fonts are available on the printer. To do so, enter this information in the font list for each printer.

Maintaining PostScript Filters

PostScript filters provided with this release are pre-installed. This section contains specific information about the location and function of these filters. In certain circumstances, you may find it helpful to change existing filter descriptions.

PostScript filters are contained in the directory **/usr/lib/lp/filter/postscript.**

There are two types of filters: fast filters and slow filters. For a definition of these types, see the **lpfilter**(1M) manual page and "Defining a Filter" earlier in this section.

A prerequisite of communication between any system and a serial PostScript printer is the presence of the **postio** filter on the system. This program is the only mandatory PostScript filter for serial PostScript printers. The following filters allow other types of documents to be translated to PostScript and to be printed on a PostScript printer.

File Content Type	Filter
simple	**postprint**
troff	**dpost**
daisy (as for Diablo 630)	**postdaisy**

dmd (as for AT&T 630)	**postdmd**
tek4014 (Tektronix)	**posttek**
plot	**postplot**

The following filters perform special functions:

Function	Filter
Communicate with printer	**postio**
Download fonts	**download**
Reverse or select pages	**postreverse**
Matrix gray scales	**postmd**

Installing and Maintaining PostScript Fonts

One of the advantages of PostScript is its ability to manage fonts. Fonts are stored in outline form, either on the printer or on a host that communicates with a printer. When a document is printed, the PostScript interpreter generates each character as needed (in the appropriate size) from the outline description of it. If a font required for a document is not stored on the printer being used, it must be transmitted to that printer before the document can be printed. This transmission process is called "downloading fonts."

Fonts are stored and accessed in several ways.

- Fonts may be stored permanently on a printer. These "printer resident" fonts may be installed in ROM on the printer by the manufacturer. If the printer has a disk, fonts may be installed on that disk by you (that is, by the print service administrator). Most PostScript printers are shipped with thirty-five standard fonts.

- A font may be "permanently-downloaded" by being transmitted to a printer with a PostScript "exitserver" program. A font downloaded in this way will remain in the printer's memory until the printer is turned off. Memory allocated to this font will reduce the memory available for PostScript print requests. Use of exitserver programs requires the printer system password and may be reserved for the printer administrator. This method is useful when there is continual use of a font by the majority of print requests serviced by that printer.

- Fonts may be prefixed to a user's print request by the user, and be transmitted as part of the user's print request. When the user's document has been printed, the space allocated to the font is freed for other print requests. The font is stored in the user's directory. This method is preferable for fonts with more limited usage.

- Fonts may be stored on a system shared by many users. These fonts may be described as "host-resident." This system may be a server for the printer or may be a system connected to the printer by a network. Each user may request fonts in the document to be printed. This method is useful when there are a large

number of available fonts or there is not continual use of these fonts by all print requests. If the fonts will be used only on printers attached to a server, they should be stored on the server. If the fonts are to be used by users on one system, who may send jobs to multiple printers on a network, they may be stored on the users' system.

The LP print service allows you to manage fonts in any of three ways. It provides a special download filter to manage fonts using the last method described above.

Managing Printer-Resident Fonts

Most PostScript printers come equipped with fonts resident in the printer ROM. Some printers have a disk on which additional fonts are stored. When a printer is installed, the list of printer-resident fonts should be added to the font-list for that printer. These lists are kept in the printer administration directories. For a particular printer, this list is contained in the file

/etc/lp/printers/*printer-name***/residentfonts**

where *printer–name* is the name of the printer.

When fonts are permanently downloaded to the printer, the font names should be added to this file. Host-resident fonts must reside on the host to which the printer is attached. Using System V Release 4 **lpNet** printing, fonts are not transferred from a local printer to the remote host. This list should include fonts which reside on that system and are available for downloading to the printer. These files must be edited manually; that is, with the help of a text editor such as **vi**(1).

For example, the **residentfonts** file for the DG Model 6771 PostScript printer with the 35 font option would include:

```
AvantGarde-Book
AvantGarde-BookOblique
AvantGarde-Demi
AvantGarde-DemiOblique
Bookman-Demi
Bookman-DemiItalic
Bookman-Light
Bookman-LightItalic
Courier
Courier-Bold
Courier-BoldOblique
Courier-Oblique
Helvetica
Helvetica-Bold
Helvetica-BoldOblique
Helvetica-Narrow
Helvetica-Narrow-Bold
Helvetica-Narrow-Bold
Helvetica-Narrow-BoldOblique
Helvetica-Narrow-Oblique
Helvetica-Oblique
NewCenturySchlbk-Bold
```

```
NewCenturySchlbk-BoldItalic
NewCenturySchlbk-Italic
NewCenturySchlbk-Roman
Palatino-Bold
Palatino-BoldItalic
Palatino-Italic
Palatino-Roman
Symbol
Times-Bold
Times-BoldItalic
Times-Italic
Times-Roman
ZapfChancery-MediumItalic
ZapfDingbats
```

Installing and Maintaining Host-Resident Fonts

Some fonts will be resident on the host and transmitted to the printer as needed for particular print requests. As the administrator, it's your job to make PostScript fonts available to all the users on a system. To do so, you must know how and where to install these fonts, using the guidelines described previously. Because fonts are requested by name and stored in files, LP keeps a map file with the correspondence between font names and the names of the files that contain the fonts. Both of these must be updated when fonts are installed on the host.

Install host-resident PostScript fonts by doing the following:

- Copy the font file to the appropriate directory.
- Add to the map table the name of the font and the name of the file in which it resides.

Where Are Fonts Stored?

The fonts available for use with PostScript printers reside in directories called **/usr/share/lib/hostfontdir/***typeface/font*. Replace *typeface* with a typeface name such as palatino or helvetica. Replace *font* with bold, italic, and so forth.

Adding an Entry to the Map Table

Also within the **hostfontdir** directory, you must create and maintain a map table file, called **map**, which shows the correspondence between the name assigned to each font by the foundry (the company that created the font) and the name of the file in which that font resides. For example, to map the font called "XYZ–SemiBold," add the following line to the map table file:

```
XYZ-SemiBold /usr/share/lib/hostfontdir/xyz/XYZ-SemiBold
```

The font file, XYZ–SemiBold, must be stored in the above directory. The font vendor will provide the suitable PostScript program for storing the font into a font dictionary. An example entry would be:

```
%!
serverdict begin 0 exitserver
14 dict begin
```

```
/FontInfo 10 dict dup begin
/Notice (Copyright XYZ Corporation 1992) \
All rights reserved.  This product is licensed, not sold, \
and may not be reproduced with out the consent of \
XYZ Corporation.) readonly dev
     /FullName (Xyz Semi Bold)def
     /FamilyName(Xyz)def
     /Weight(Semi Bold (Demi))def
     /version(1.1)def
     /ItalicAngle 0 def
     /isFixedPitch false def
     /UnderlinePostion -175 def
     /UnderlineThickness 41 def
end readonly def
/FontName /XYZ-SemiBold def
/Encoding StandardEncoding def
/PaintType 0 def
/FontType 1 def
/FontMatrix [0.001 0 0 0.001 0 0] readonly def
/FontBox {-155 -214 1006 865} readonly def
/Unique ID nnnnnnnnn def
currentdict end
currentfile eexect
     font vector description
cleartomark
```

Once this entry exists in the map table file on your system, your users will be able to have a XYZ–SemiBold font used in their print jobs. When they submit a file containing a request for this font, the LP print service will prepend a copy of the file

```
/usr/share/lib/hostfontdir/xyz/XYZ-SemiBold
```

to the user's file before sending it to the printer.

Downloading Host-Resident Fonts

The creators of the PostScript language anticipated that users would want to download fonts to printers. For this purpose, they defined a standard set of *structuring conventions* for PostScript programs. The download filter relies on these structuring conventions to determine which fonts must be downloaded. See your PostScript documentation for more information.

When the LP print service receives a request for a job that requires fonts not loaded on the printer, it forwards that request to a filter that downloads fonts.

The download filter does five things:

- It searches the PostScript document to determine which fonts have been requested. These requests are documented with the following PostScript structuring comments:

  `%%DocumentFonts:` *`font1 font2 ...`*

  in the header comments.

- It searches the list of fonts resident on that printer to see if the requested font must be downloaded.
- If the font is not resident on the printer, it searches the host-resident font directory (by getting the appropriate file name from the map table) to see if the requested font is available.
- If the font is available, the filter takes the file for that font and prepends it to the file to be printed.
- The filter sends the font definition file and the source file (the file to be printed) to the PostScript printer.

For example, a document may have the following font request:

```
%%Documentfonts: Helvetica XYZ-SemiBold
```

In our host-resident list above, the font, XYZ–SemiBold, is not listed. However being listed in the map table, XYZ–SemiBold is available for downloading.

Permanently Downloaded Fonts

If you frequently request a common set of fonts, you might wish to permanently download them to the printer. Note that these fonts will decrease the amount of memory available on the printer and must be downloaded each time the power is cycled on the printer.

For example, follow these steps to permanently download the font, ZZZ–Bold:

1. Make sure an entry in the resident font file, **/etc/lp/printers/printer–name/residentfonts** exists for ZZZ–Bold. This entry prevents the font being downloaded for each print request.
2. Initialize the printer by issuing the command

   lp /usr/share/lib/hostfontdir/zzz/ZZZ-Bold ↵

   to initially download the font. The font file should contain PostScript commands similar to the following. This header should be supplied by the font vendor.

```
%!
serverdict begin 0 exitserver
14 dict begin
/FontInfo 10 dict dup begin
     /Notice (Copyright ZZZ Corporation 1992) \
All rights reserved.  This product is licensed, not sold, \
and may not be reproduced with out the consent of \
ZZZ Corporation.) readonly dev
     /FullName (Zzz Semi Bold)def
     /FamilyName(Zzz)def
     /Weight(Semi Bold (Demi))def
     /version(1.1)def
     /ItalicAngle 0 def
     /isFixedPitch false def
```

```
     /UnderlinePostion -175 def
     /UnderlineThickness 41 def
end readonly def
/FontName /ZZZ-SemiBold def
/Encoding StandardEncoding def
/PaintType 0 def
/FontType 1 def
/FontMatrix [0.001 0 0 0.001 0 0] readonly def
/FontBox {-155 -214 1006 865} readonly def
/Unique ID nnnnnnnnn def
currentdict end
currentfile eexec
     font vector description
cleartomark
```

Customizing the Print Service

Although the LP print service has been designed to be flexible enough to handle most printers and printing needs, it doesn't handle every possible situation. You may buy a printer that doesn't quite fit into the way the LP print service handles printers, or you may have a printing need that the standard features of the LP print service don't accommodate.

You can customize the LP print service in a few ways. This section tells you how you can do the following:

- Adjust the printer port characteristics
- Adjust the **terminfo** database
- Write an interface program
- Write a filter
- Do character set translations

The diagram in Figure 12–5 gives an overview of the processing of a print request.

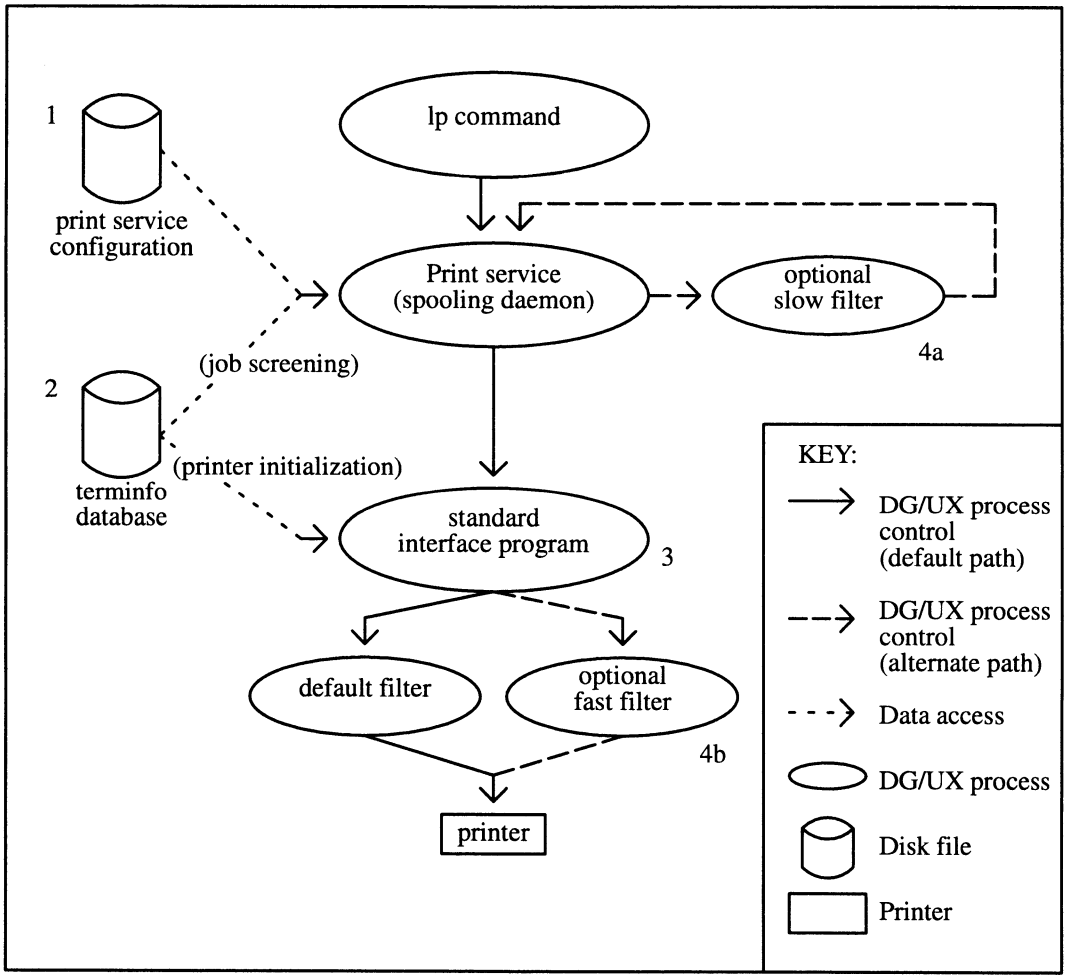

Figure 12–5 How LP Processes Print a Request

Each print request is sent to a spooling daemon that keeps track of all requests. The daemon, which is created when you start the LP print service, is also responsible for keeping track of the status of printers and slow filters; when a printer finishes printing a user's file, the daemon starts it printing another request (if there is one queued).

To customize the print service, adjust or replace some of the pieces shown in Figure 12–5. (The numbers are keyed to the diagram.)

1. For most printers, you need only change the printer configuration stored on disk. The earlier sections of this chapter explain how to do this. Configuration data that are relatively dependent on the printer include the printer port characteristics: baud rate, parity, and so on.
2. For a printer that is not represented in the **terminfo** database, you can add a new entry that describes its capabilities. The **terminfo** database is used in two parallel capacities: screening print requests to ensure that those accepted can be handled by the desired printer, and setting the printer in a state where it is ready to print a request.

   For instance, if the **terminfo** database does not contain an entry for a printer capable of setting a page length requested by a user, the spooling daemon will

reject the request. On the other hand, if it does contain an entry for such a printer, then the same information will be used by the interface program to initialize the printer.

3. For particularly difficult printers, or if you want to add features not provided by the delivered LP print service, you can change the standard interface program. This program is responsible for managing the printer: it prints the banner page, initializes the printer, and invokes a filter to send copies of a user's files to the printer.

4a. and 4b.
To provide a link between the applications used on your system and the printers, you can add slow and fast filters. Each type of filter can convert a file into another form, mapping one set of escape sequences into another, for instance, and can provide special setup by interpreting print modes requested by a user. Slow filters are run separately by the daemon, to avoid tying up a printer. Fast filters are run so their output goes directly to the printer; thus they can exert control over the printer.

Adjusting the Printer Port Characteristics

You should make sure that the printer port characteristics set by the LP print service match the printer communication settings. The standard printer port settings have been designed to work with typical files and many printers, but they won't work with all files and printers. This isn't really a customizing step, because a standard feature of the LP print service is to allow you to specify the port settings for each printer. However, it's an important step in getting your printer to work with the LP print service, so it's described in more detail here.

When you add a new printer, read the documentation that comes with it so that you understand what it expects from the host (the LP print service). Then read the manual page for **stty**(1). It summarizes the various characteristics that can be set on a terminal or printer port.

Only some of the characteristics listed in the **stty**(1) manual page are important for printers. The ones likely to be of interest to you appear in Table 12–5 (but you should still consult the **stty**(1) manual page for others).

Table 12–5 stty Options Related to Printers

stty Option	Meaning
evenp	Send even parity in the 8th bit
oddp	Send odd parity in the 8th bit
-parity	Do not generate parity; send all 8 bits unchanged
110 - 38400	Set the communications speed to this baud rate
ixon	Enable XON/XOFF (also known as START/STOP or DC1/DC3) flow control
-ixon	Turn off XON/XOFF flow control

Continued

Table 12–5 stty Options Related to Printers

stty Option	Meaning
–opost	Do not do any "output post-processing"
opost	Do "output post-processing" according to the settings listed below
onlcr	Send a carriage return before every linefeed
–onlcr	Do not send a carriage return before every linefeed
ocrnl	Change carriage returns into linefeeds
–ocrnl	Do not change carriage returns into linefeeds
–tabs	Change tabs into an equivalent number of spaces
tabs	Do not change tabs into spaces

When you have a set of printer port characteristics you think should apply, adjust the printer configuration as described in "Printer Port Characteristics" under "Configuring Your Printers" in this section. You may find that the default settings are sufficient for your printer.

Adjusting the terminfo Database

The LP print service relies on a standard interface and the **terminfo** database to initialize each printer and establish a selected page size, character pitch, line pitch, and character set. Thus, it is usually sufficient to have the correct entry in the **terminfo** database to add a new printer to the LP print service. Several entries for Data General printers and other popular printers are delivered in the standard **terminfo** database.

Each printer is identified in the **terminfo** database with a short name; this kind of name is identical to the kind of name used to set the **TERM** shell variable. The complete list of supported terminal and printer names are in the **terminfo** database. To view the entire list, use this command line:

```
# ls -CR  /usr/share/lib/terminfo/*  |  more  ↵
```

If you cannot find a **terminfo** entry for your printer, you should add one. If you do not, you may still be able to use the printer with the LP print service, but you will not have the option of automatic selection of page size, pitch, and character sets, and you may have trouble keeping the printer set in the correct modes for each print request. Another option to follow, instead of updating the **terminfo** entry, is to customize the interface program used with the printer. (See the next section for details on how to do this.)

There are hundreds of items that can be defined for each terminal or printer in the **terminfo** database. However, the LP print service uses fewer than 50 of these. Table 12–6 lists some of the items that need to be defined (as appropriate for the printer) to add a new printer to the LP print service.

Table 12–6 terminfo Items Relevant to Printers

terminfo Item	Meaning
Booleans:	
cpix	Changing character pitch changes resolution
daisy	Printer needs operator to change character set
lpix	Changing line pitch changes resolution
Numbers:	
bufsz	Number of bytes buffered before printing
cols	Number of columns in a line
cps	Average print rate in characters per second
it	Tabs initially every # spaces
lines	Number of lines on a page
orc	Horizontal resolution in units per character
orhi	Horizontal resolution in units per inch
orl	Vertical resolution in units per line
orvi	Vertical resolution in units per inch
Strings:	
chr	Change horizontal resolution
cpi	Change number of characters per inch
cr	Carriage return
csnm	List of character set names
cud1	Down one line
cud	Move carriage down # lines
cuf	Move carriage right # columns
cuf1	Carriage right
cvr	Change vertical resolution
ff	Page eject
hpa	Horizontal position absolute
ht	Tab to next 8-space tab stop
if	Name of initialization file
iprog	Pathname of initializing program
is1	Printer initialization string
is2	Printer initialization string
is3	Printer initialization string

Continued

Table 12–6 terminfo Items Relevant to Printers

terminfo Item	Meaning
lpi	Change number of lines per inch
mgc	Clear all margins (top, bottom, and sides)
rep	Repeat a character # times
rwidm	Disable double wide printing
scs	Select character set
scsd	Start definition of a character set
slines	Set page length to # lines
smgb	Set bottom margin at current line
smgbp	Set bottom margin
smgl	Set left margin at current column
smglp	Set left margin
smgr	Set right margin at current column
smgrp	Set right margin
smgt	Set top margin at current line
smgtp	Set top margin
swidm	Enable double wide printing
vpa	Vertical position absolute

To construct a database entry for a new printer, see details about the structure of the **terminfo** database in the **terminfo**(4) manual page.

Once you've made the new entry, you need to compile it into the database using the **tic**(1) program. Just enter the following command:

```
# tic file-name ↵
```

file–name is the name of the file containing the **terminfo** entry you have crafted for the new printer. The LP print service gains much efficiency by caching information from the **terminfo** database. If you add or delete **terminfo** entries, or change the values that govern pitch settings, page width and length, or character sets, you should stop and restart the LP print service so it can read the new information.

How to Write an Interface Program

If you have an interface program that you used with the LP service before DG/UX Release 5.4, it should still work with the LP print service. Note, though, that several **–o** options have been "standardized," and will be passed to every interface program. These may interfere with similarly named options used by your interface.

If you have a printer that is not supported by simply adding an entry to the **terminfo** database, or if you have printing needs that are not supported by the

standard interface program, you can furnish your own interface program. It is a good idea to start with the standard interface program, and change it to fit, rather than starting from scratch. You can find a copy of it under the name **/var/spool/lp/model/standard**.

What Does an Interface Program Do?

Any interface program is responsible for doing the following tasks:

- Initializing the printer port, if necessary. The generic interface program uses the **stty** command to do this.
- Initializing the physical printer. The generic interface program uses the **terminfo** database and the **TERM** shell variable to get the control sequences to do this.
- Printing a banner page, if necessary.
- Printing the correct number of copies of the request content.

An interface program is not responsible for opening the printer port. The LP print service opens the port, a process which includes calling a "dial-up" printer, if one is used to connect the printer. The printer port connection is given to the interface program as standard output, and the printer is identified as the "controlling terminal" for the interface program so that a "hang-up" of the port will cause a **SIGHUP** signal to be sent to the interface program.

A customized interface program must not terminate the connection to the printer or "uninitialize" the printer in any way.

How Is an Interface Program Used?

When the LP print service routes an output request to a printer, the interface program for the printer is invoked as follows:

/var/spool/lp/admins/lp/interfaces/*P ID user title copies options f1 f2 ...*

Arguments for the interface program are:

P	printer name
id	request ID returned by the **lp**(1) command
user	username of the user who made the request
title	optional title specified by the user
copies	number of copies requested by the user
options	blank-separated list of options specified by the user or set by the LP print service
f1, f2, ...	full pathnames of the files to be printed

When the interface program is invoked, its standard input comes from **/dev/null**, its standard output is directed to the printer port, and its standard error output is directed to a file that will be given to the user who submitted the print request.

The standard interface recognizes the following values in the blank-separated list in `options`.

nobanner This option is used to skip the printing of a banner page; without it, a banner page is printed.

nofilebreak This option is used to skip page breaks between separate data files; without it, a page break is made between each file in the content of a print request.

cpi=*decimal–number1*
lpi=*decimal–number2*

These options specify a format of *decimal–number1* columns per inch and *decimal–number2* lines per inch, respectively. The standard interface program extracts from the **terminfo** database the control sequences needed to initialize the printer to handle the character and line pitches.

The words **pica**, **elite**, and **compressed** are acceptable replacements for *decimal–number1* and are synonyms, respectively, for **10** columns per inch, **12** columns per inch, and as many columns per inch as possible.

length=*decimal–number1*
width=*decimal–number2*

These options specify the length and width, respectively, of the pages to be printed. The standard interface program extracts from the **terminfo** database the control sequences needed to initialize the printer to handle the page length and page width.

stty='*stty–option–list***'**

The *stty–option–list* is applied after a default *stty–option–list* as a set of arguments to the **stty** command. The default list is used to establish a default port configuration; the additional list given to the interface program is used to change the configuration as needed.

lpd='*argument–list***'**

This option is used internally by the **lpsched** command; you can ignore it.

flist='*file–list***'** This option is used internally by the **lpsched** command; you can ignore it.

The above options may be specified by the user when issuing a print request. Alternatively, they may be specified by the LP print service from defaults given by the administrator either for the printer (**cpi**, **lpi**, **length**, **width**, **stty**) or for the preprinted form used in the request (**cpi**, **lpi**, **length**, **width**).

Additional printer configuration information is passed to the interface program in the following shell variables:

TERM=*printer–type*

This shell variable specifies the type of printer. The value is used as a key for getting printer capability information from the **terminfo** database.

FILTER='*pipeline***'**
This shell variable specifies the filter to use to send the request content to the printer; the filter is given control of the printer.

CHARSET=*character–set*
This shell variable specifies the character set to be used when printing the content of a print request. The standard interface program extracts from the **terminfo** database the control sequences needed to select the character set.

A customized interface program should either ignore these options and shell variables or should recognize them and treat them in a consistent manner.

Customizing the Interface Program

Make sure that the custom interface program sets the proper **stty** modes (terminal characteristics such as baud rate and output options). The standard interface program does this, and you can follow suit. Look for the section that begins with the shell comment

```
## Initialize the printer port
```

Follow the code used in the standard interface program. It sets both the default modes and the adjusted modes given by either the LP print service or the user with a line such as the following:

```
stty mode options 0<&1
```

This command line takes the standard input for the **stty** command from the printer port. An example of an **stty** command line that sets the baud rate at 1200 and sets some of the option modes is shown below.

```
stty -parenb -parodd 1200 cs8 cread clocal ixon 0<&1
```

One printer port characteristic not set by the standard interface program is hardware flow control. The way that this is set will vary, depending on your computer hardware. The code for the standard interface program suggests where this and other printer port characteristics can be set. Look for the section that begins with the shell comment

```
# Here you may want to add other port initialization code.
```

Because different printers have different numbers of columns, make sure the header and trailer for your interface program correspond to your printer. The standard interface program prints a banner that fits on an 80-column page (except for the user's title, which may be longer). Look in the code for the standard interface program for the section that begins with the shell comment

```
## Print the banner page
```

The custom interface program should print all user related error messages on the standard output or on the standard error. The messages sent to the standard error will be mailed to the user; the messages printed on the standard output will end up

on the printed page where they can be read by the user when he or she picks up the output.

When printing is complete, your interface program should exit with a code that shows the status of the print job. Exit codes are interpreted by the LP print service as shown in Table 12–7.

Table 12–7 Printer Exit Codes

Code	Meaning to the LP print service
0	The print request has been completed successfully. If a printer fault has occurred, it has been cleared.
1 to 127	A problem has been encountered in printing this particular request (for example, too many non-printable characters, or the request exceeds the printer capabilities). The LP print service notifies the person who submitted the request that there was an error in printing it. This problem will not affect future print requests. If a printer fault had occurred, it has been cleared.
128	Reserved for internal use by the LP print service. Interface programs must not exit with this code.
129	A printer fault has been encountered in printing the request. This problem will affect future print requests. If the fault recovery for the printer directs the LP print service to wait for the administrator to fix the problem, the LP print service will disable the printer. If the fault recovery is to continue printing, the LP print service will not disable the printer, but will try printing again in a few minutes.
greater than 129	These codes are reserved for internal use by the LP print service. Interface programs must not exit with codes in this range.

As the table shows, one way of alerting the administrator to a printer fault is to exit with a code of 129. Unfortunately, if the interface program exits, the LP print service has no choice but to reprint the request from the beginning when the fault has been cleared. Another way of getting an alert to the administrator (that does not require the entire request to be reprinted) is to have the interface program send a fault message to the LP print service but wait for the fault to clear. When the fault clears, the interface program can resume printing the user's file. When the printing is finished, the interface program can give a zero exit code just as if the fault had never occurred. An added advantage is that the interface program can detect when the fault is cleared automatically, so that the administrator doesn't have to enable the printer.

Fault messages can be sent to the LP print service using the **lp.tell** program. This is referenced using the LPTELL shell variable in the standard interface code. The program takes its standard input and sends it to the LP print service where it is put into the message that alerts the administrator to the printer fault. If its standard input is empty, **lp.tell** does not initiate an alert. Examine the standard interface code immediately after these comments for an example of how the **lp.tell** (LPTELL environment variable) program is used:

```
# Here's where we set up the $LPTELL program to capture
# fault messages.

# Here's where we print the file.
```

If the special exit code 129 or the **lp.tell** program is used, there is no longer a need for the interface program to disable the printer itself. Your interface program can disable the printer directly, but doing so will override the fault alerting mechanism. Alerts are sent only if the LP print service detects the printer has faulted, and the special exit code and the **lp.tell** program are its main detection tools.

If the LP print service has to interrupt the printing of a file at any time, it will terminate the interface program with a signal `TERM` (signal 15; see **kill**(1) and **signal**(2)). If the interface program dies from receipt of any other signal, the LP print service assumes that future print requests won't be affected, and continues to use the printer. The LP print service notifies the person who submitted the request that the request has not been finished successfully.

When the interface is first invoked, the signals `HUP`, `INT`, `QUIT`, and `PIPE` (trap numbers 1, 2, 3, and 13) are ignored. The standard interface changes this so that these signals are trapped at appropriate times. The standard interface interprets receipt of these signals as warnings that the printer has a problem; when it receives one, it issues a fault alert.

How to Write a Filter

A filter is used by the LP print service each time it has to print a type of file that isn't acceptable by a printer. A filter can be as simple or as complex as needed; there are only a few external requirements:

- The filter should get the content of a user's file from its standard input and send the converted file to the standard output.
- A **slow** filter can send messages about errors in the file to standard error. A **fast** filter should not, as described below. Error messages from a **slow** filter are collected and sent to the user who submitted the file for printing.
- If a **slow** filter dies because of receiving a signal, the print request is stopped and the user who submitted the request is notified. Likewise, if a **slow** filter exits with a non-zero exit code, the print request is stopped and the user is notified. The exit codes from **fast** filters are treated differently, as described below.
- A filter should not depend on other files that normally would not be accessible to a regular user; if a filter fails when run directly by a user, it will fail when run by the LP print service.

The "Providing Filters" section earlier in this section describes how to add a filter to the LP print service.

If you want your filter to detect printer faults, you must also fulfill the following requirements:

- If possible, the filter should wait for a fault to be cleared before exiting. Additionally, it should continue printing at the top of the page where printing stopped after the fault clears. If the administrator does not want this contingency followed, the LP print service will stop the filter before alerting the administrator.

- It should send printer fault messages to its standard error as soon as the fault is recognized. It does not have to exit, but can wait as described above.

- It should not send messages about errors in the file to standard error. These should be included in the standard output stream, where they can be read by the user.

- It should exit with a zero exit code if the user's file is finished (even if errors in the file have prevented it from being printed correctly).

- It should exit with a non-zero exit code only if a printer fault has prevented it from finishing a file.

- When added to the filter table, it must be added as a **fast** filter. (See "Defining a Filter" in this chapter for details.)

Translating Character Sets

To print a file that uses a character set different from the one used by your printer, use the **iconv** command to convert the file's character set to one that is compatible with the printer's. The **iconv** command uses the database, **/usr/lib/kbd/iconv_data**, which contains a list of code sets that can be converted. The code set being converted must be listed in this file. See the **iconv**(1) manual page for more information on this command and the **iconv_data** file.

For example, suppose you wanted to print a file that was created on an AOS/VS system using 8 bit DGI characters on a printer that supports IBM character sets 437 and 850. You can convert both of these files with standard System V Release 4 printer filters using **iconv** as the filter. Follow these steps to translate the files:

1. Create a filter with the input defined as "dgi" and output type of "88591". The only supported conversion tables for DGI are to and from 88591. So, you must first map all DGI characters to ISO 8859.1 characters first and then map again for the required device character set. In the filter, you specify the command **iconv** as the program, the from (**-f**) value as dgi, and the to (**-t**) value as 88591. These values, are taken from the **iconv_data** file. The from value is column 1 in the file. The to value is column 2.

   The entries for input, output, and content types are not specific indicators for the filter **iconv**. However, the values must match between the filters and the printers. The specific conversion is specified on the filter command line, so the entry, "dgi", is the descriptive reference for the input type to the filter.

   Here is the specific example for a filter called "dgi_to_88591":

```
Input types: dgi
Output types: 88591
Printer types: any
Printers: any
Filter type: slow
Command: /usr/bin/iconv -f dgi -t 88591 -m b
```

2. Define a second filter to convert from the ISO 8859.1 character set to the IBM code page 850. Note that all the DGI characters map to the ISO 8859.1 character set, but not all are supported in the IBM 850 character set.

   Define the input type for the filter as 88591 and the output type as PC850. **iconv** is again specified as the program with the from (**–f**) value as 88591 and the to (**–t**) value as PC850.

   Here is the specific example for a filter called **88591_to_PC850**:

```
Input types: 88591
Output types: PC850
Printer types: any
Printers: any
Filter type: slow
Command: /usr/bin/iconv -f 88591 -t PC850 -m b
```

3. Define a printer of type "ibmgraphics." The Content type is PC850, which matches the output type of the filter, 88591_to_PC850. Below is the specific example for the printer called ibm_850. On the ibmgraphics printer, character_set_1 is code page 437, and character_set_2 is code page 850.

```
printer ibm_850 is idle. enabled since Mon Jul 13 16:56:13 EDT 1992.
available.
        Form mounted:
        Content types: PC850
        Printer types: ibmgraphics
        Description: direct connect ibm proprinter
        Connection: direct
        Interface: /usr/lib/lp/model/standard
        On fault: mail to transou once
        After fault: continue
        Users allowed:
                (all)
        Forms allowed:
                (none)
        Banner not required
        Character sets:
                character_set_1
                character_set_2
        Default pitch: 10 CPI 6 LPI
        Default page size: 80 wide 66 long
        Default port settings:
```

To print a file that contains dgi characters on an IBM ProPrinter which supports the code page 850, you would enter the following **lp** command.

```
lp -d ibm_850 -T dgi -S character_set_2 ./dgi_file ↵
```

In this command line **–d** selects the printer **ibm_850**, **–T** specifies the content type of the file as **dgi**, and **–S** selects character set 2 (code page 850) for the IBM printer.

The entire conversion of file follows this path:

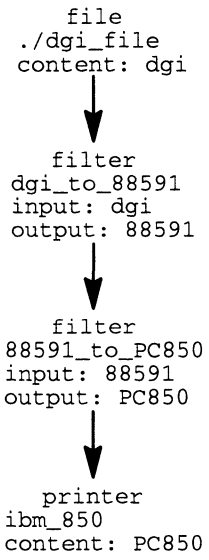

As another example, to print a file containing characters from the ISO 8859.1 character set on an IBM ProPrinter supporting the code page 850, you would enter the following lp command.

```
lp -d ibm_850 -T 88591 -S character_set_2 ./88591_file ↵
```

In this command line, **–d** selects the printer ibm_850, **–T** specifies the content type of the file is 88591, and **–S** selects character set 2 (code page 850) for the IBM printer.

The file conversion follows this path:

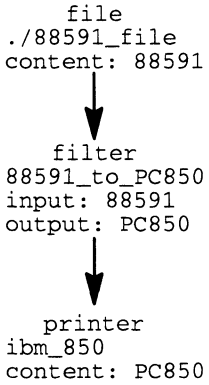

Using the man pages for **iconv**(1), **lp**(1), **lpadmin**(1), and these examples, you should be able to define the necessary conversion filters for any of the characters sets supported in **/usr/lib/kbd** and the map file **/usr/lib/kbd/iconv_data**.

Command Reference for LP Print Service Administration

The commands in Table 12–8 are found in the **/usr/bin** directory unless otherwise noted. To use the administrative commands, you must be the superuser (**sysadm** or **root**).

Table 12–8 LP Print Service Command Reference

Action	Command
Activating a printer:	**enable**(1)
Canceling a request for a file to be printed:	**cancel**(1)
Sending a file (or files) to a printer:	**lp**(1)
Reporting the status of the LP print service:	**lpstat**(1)
Deactivating specified printers:	**disable**(1)
Permitting job requests to be queued for a specific destination:	**accept**(1M)
Preventing jobs from being queued for a specified destination:	**reject**(1M)
Setting up or changing printer configurations:	**lpadmin**(1M)
Setting up or changing filter definitions:	**lpfilter**(1M)
Setting up or changing preprinted forms:	**lpforms**(1M)
Mounting a form:	**lpadmin**(1M)
Moving output requests from one destination to another:	**lpmove**(1M)
Starting the LP print service scheduler:	**lpsched**(1M)
Stop the LP print service scheduler:	**lpshut**(1M)
Setting or changing the default priority and priority limits that can be requested by users of the LP print service:	**lpusers**(1M)

End of Chapter

Chapter 13
Login Account Management

This chapter is for systems supporting one or more login accounts. As a system manager, you will set up and manage the everyday working environment in which users function. This task includes adding and removing login accounts, creating new aliases and groups, answering users' questions, and helping users with system problems. If your system is part of a network that uses the Network Information Service (NIS) facility, formerly called the Yellow Pages (YP), some of these tasks will be the responsibility of the master server administrator. For more information on the NIS facility, see *Managing ONC™/NFS® and Its Facilities on the DG/UX™System.*

Login Account Terms

Read the following definitions before beginning the procedures in this chapter.

login name A valid name for logging on to the system. Also known as a *username*. A login name can be up to 32 alphabetic or numeric characters, but if you want to be compatible with traditional systems, the login name should be no more than eight characters long. Note that by default only the first eight characters of a login name are checked for uniqueness. See the **passwd**(4) manual page for more information.

password A unique string of alphabetic or numeric characters that allows a user access to the system. A password must be at least six characters, and a maximum of eight characters. At least one character *must* be a numeral or a special character. If you do not set an initial password when adding a user with the Add operation, the user will be unable to use the account until you do set an initial password.

password aging A system that forces users to change their passwords periodically. An "aged" password is one that must be changed within the specified number of weeks. When this time period lapses, the password will no longer gain entry to the system. The user must choose a new password. System administrators decide if they want to use password aging or not. See "Expert Information" later in this chapter for details.

group A set of users with common access privileges to a common set of files based on group ID numbers. Also known as a *group name*. People that need access to the same files can be listed in a group. For example, anyone in group **prog** could access those files associated

with that group name, assuming the group permissions were set correctly for the files. Users not in the group **prog** do not have the same access rights to the files as do members of **prog**.

alias A mailing list. An alias contains one or more user login names or aliases. If you address mail to an alias, the mail is delivered to all users listed in that alias. An alias entry consists of an alias name and a list of alias members.

user ID A unique number that identifies a user to the system. The user ID number (**uid**) is between 100 and 60,000. The number must not include a comma. Superuser (**sysadm** and **root**) use 0. System ID numbers are 1 to 99.

group ID A unique number that identifies a group to the system. The same conditions apply to the group ID (**gid**) number as to the **uid**. System ID numbers are 1 to 99.

home directory
The origination point of the user's directory tree. The home directory is where the user is placed upon logging in. The name is usually the same as the login name, preceded by a parent directory. For example, user **poulet**'s home directory might be **/mach_2/poulet**.

initial program
The program invoked at the time the user logs in. Choices are:

/sbin/sh The Bourne shell. See **sh**(1).

/usr/bin/csh The C shell. See **csh**(1).

/usr/bin/ksh The Korn shell. See **ksh**(1).

You may specify any executable that you or the user prefers.

NIS master server
The NIS master server is the single system in the network that holds the master networks and hosts files that are exported to other systems. Global changes can be made only on the master system. If your system has the Network File System (ONC/NFS), and the current system is the NIS master server, you will be asked to choose local or global options in queries for user management in this chapter.

About Login Accounts

The operation for adding user login accounts has preset defaults that may be adequate for your environment. You can, if you wish, tailor these defaults using the **sysadm** operation `User -> Login Account -> Defaults -> Set`.

The following list shows how defaults are originally set on the DG/UX system:

User Id Set to highest unassigned valid number plus one, as long as it is less than 60,000; otherwise, lowest unassigned number over 100.

Group Set to **general**.

Home directory Set to **/home/***login_name*.

User comment No default.

Shell program Set to **/sbin/sh**.

When you are ready, read "The User's Environment" later in this chapter. In it, we present information on global and local user profiles, environment variables, file creation permissions, default and restricted shells, an electronic bulletin board, and system mail. In addition, we offer suggestions for tracking user problems and provide a sample Trouble Log for users to fill out.

Using Groups and Aliases

Let's say that users on your system are divided into two categories: programmers and data entry people. Initially, you can assign everybody to the default group that comes with the DG/UX system. Members of group **general** are allowed access to all directories and files owned by **general**. This is a shared ownership. Later, you can assign users to additional groups. For instance, you might put programmers in group **prog** and data entry people in group **pool**. Note that group **general** is simply provided as a default to speed up the process of adding users. You can rename it or delete it. Later, you can create aliases and groups based on tasks, projects, or whatever you choose.

Aliases are simply mailing lists used by the **mailx**(1) command. The default is **everybody**. See "Adding Mail Aliases," later in this chapter, for details.

Specifying a Parent Directory and Initial Program

Although **/home/***login_name* is the default home directory, we can't be sure of how you've laid out your virtual disks and file systems when you installed the DG/UX system. You will have to indicate the name of the directory that you wish to make the default parent directory for users.

If you have or plan to have ONC/NFS in the future, you should make sure that parent directories have unique names across systems. One way to do this is by using the host name as the parent directory name. This practice makes it easy to identify the origin and ownership of file systems in your network.

The default initial program is traditionally specified as **/sbin/sh**, but it can be any executable local shell program.

Login Account Procedures

The User menu of **sysadm** provides menus whose operations let you manage user login accounts and user login account defaults, user groups, and user mail aliases.

The following sections discuss the menus and operations in more detail.

Managing Login Accounts

A login account determines some of the characteristics of a user's environment as well as some attributes of programs the user runs. Local login account information is in the file **/etc/passwd**. Each line in the file represents a user's login account.

The Login Account menu provides operations for adding, deleting, modifying, and listing login accounts. There is also a submenu for managing the default values that appear in the Add operation when you add a login account.

Adding Login Accounts

Select the Add operation to add a user login account. The operation adds the entry to your **/etc/passwd** file unless your system is the master server in an NIS domain, in which case you have the option of adding the entry either to the NIS **yppasswd** database or to your local **passwd** file.

For each login account you add, the operation queries you with the following prompts:

Login Name
: A login name may be up to 32 characters, but to be compatible with more traditional systems, the login name should be no more than eight characters. Also, by default only the first eight characters of a login name are checked for uniqueness, so any names longer than eight characters must not have the same first eight characters. The name must be unique on your system. If your system is in an NIS domain, the name should be unique within the domain.

Add password aging for the login account
: Select this feature to set minimum and maximum age limits on the login password. Password aging is a security feature that lets you control how long a password may be in effect before the user is forced to change it. If you select password aging, the operation presents two additional queries about password aging at the end of the operation. The queries, discussed later, are **Minimum number of weeks before** *user* **may change his password** and **Maximum number of weeks until** *user* **must change his password**.

User Id
: The user ID, or UID, is a number between 0 and 60,000. All numbers below 100 are reserved for the system accounts. The UID should be unique on your system. If your system is in an NIS domain, the UID should be unique within the domain. Other than serving as a unique identifier, the UID is arbitrary. The default is the highest assigned UID plus one, as long as the number is 60,000 or less; otherwise, the default is the lowest unassigned UID.

Group The user group must already exist. Add a user group with the operation `User -> Group -> Add`.

Home Directory

This is the directory that will be the user's home directory and current directory upon logging into the system. If the directory does not already exist, you may direct the operation to create it in a query that appears later in the operation.

User Comment

This comment may be any text that you choose to enter to describe the user or the login account.

Shell Program

Select a shell for the user. The choice is largely arbitrary and depends on the tastes of the user. For a comparison of the shells, see *Using the DG/UX™ System*.

Create home directory for *username*

If the directory you specified in the **Home Directory** prompt does not already exist, select this option to create it.

Set an initial password for *username*

Select this option if you want to assign a password for the user at this time. If you do not assign a password, the operation will assign an impossible password so that no one can log into the system through the account. To change the password for such an account after creating it with the Add operation, use the **passwd**(1) command (if the account is in the local **passwd** database) or the **yppasswd**(1) command (if the account is in the NIS **yppasswd** database). It is good practice to have the user change the password as soon as possible.

Minimum number of weeks before *user* **may change his password**

This query appears only if you selected password aging for the login account. Enter the minimum number of weeks that may pass before the user may change his password. If the user attempts to change his password before this time has passed, the attempt will fail. To let the user change his password at any time up to the expiration date, enter zero. Set the expiration date in the following query.

Maximum number of weeks until *user* **must change his password**

This query appears only if you selected password aging for the login account. Enter the maximum number of weeks that may pass before the user must change his password. If the user does not change his password by this expiration period, the user will be prompted for a new password the next time he logs into the system.

If both the minimum number of weeks, set in the previous query, and the maximum number of weeks are equal to zero, the user must change his password the first time he logs into the system. Setting the minimum and maximum both to zero forces only one change of password for the account. After that change has occurred (even if performed by the superuser), password aging does not apply and the password may remain indefinitely.

If the minimum number of weeks is greater than the maximum number of weeks, only the superuser can change the password for this account.

After you have answered these queries, the Add operation proceeds to add the **passwd** entry.

If the user's home directory already exists or if the Add operation creates it, the Add operation copies to it some initial personal configuration files from a *skeleton* directory, **/etc/skel**. This skeleton directory exists simply as a prototype of a home directory, including several basic configuration files such as **.login** and **.profile** that serve as starting points for the user to begin customizing the working environment.

You may wish to create your own skeleton directory based on the one shipped with the system. Do not add or change files in the system default skeleton directory because future revisions of the DG/UX system may overwrite it with new contents. To set a different default skeleton directory name, use the operation `User -> Login Account -> Defaults -> Set`.

Deleting Login Accounts

To remove a user login account from the password database, select the Delete operation. The operation removes the login entry from your **/etc/passwd** file unless your system is the master server in an NIS domain, in which case the operation lets you choose whether you want to remove the entry from the global NIS **yppasswd** database or from your system's local **passwd** database.

The Delete operation lets you choose whether to delete the user's home directory or not. The operation does not delete the user's mail file (**/var/mail/***username*).

If you remove a user's login account while the user is logged in, the user does not experience any interruption in service. If the user logs out, however, the user will not be able to log in again.

After you have deleted a login account, the system can no longer resolve references to the user ID number and get the login name. This means that commands such as **ls**(1) will return the user ID number instead of the name, where applicable. If you assign a new login account with an old ID number, the new login name will appear associated with all files that had been associated with the old login name. The algorithm used to determine a default UID for a new login account attempts to avoid this sort of confusion.

Modifying Login Accounts

To change a user's login account, select the Modify operation. The Modify operation presents the same queries that you see when you use the Add operation to create the user login account.

The Modify operation looks for the login account in your system's local **/etc/passwd** file unless your system is the NIS master server, in which case the operation lets you choose whether to change the local **passwd** file or the global NIS **yppasswd** file.

Changes to a local login account will be in effect when the operation completes, but changes to a global NIS login account may not be in effect until the following day.

Changes to a login name and user number take effect throughout the system as soon as the **passwd** file changes; however, other changes to the login account (such as a changed home directory) will not take effect until the user logs in again.

Displaying Login Accounts

Select the List operation to display login accounts. The operation lists information from the local **/etc/passwd** file unless you are a in an NIS domain, in which case the operation lets you choose which database to list.

You may restrict the login accounts listed by specifying a user ID, group ID, and login name. The default for all three fields is **all**. The operation lists accounts that match all three fields.

A typical listing of all accounts in a local login database looks like this:

```
Username        Uid     Gid  Home Directory            Shell
--------        ---     ---  --------------            -----
root              0       1  /                         /sbin/sh
xdm               0       1  /                         /usr/bin/X11/telxdm
sysadm            0       0  /admin                    /sbin/sh
daemon            1       1  /                         /sbin/sh
bin               2       2  /bin
sys               3       3  /usr/src
adm               4       4  /usr/adm                  /sbin/sh
uucp              5       5  /usr/spool/uucp           /usr/lib/uucp/uucico
nuucp             5       1  /usr/lib/uucp             /sbin/sh
lp                6       2  /usr/lib                  /sbin/sh
mail              8       1  /usr/mail                 /usr/bin/mail
sync             19       1  /                         /bin/sync
yp               37      37  /usr/etc/yp               /sbin/sh
nfs              38      38  /                         /sbin/sh
ftp             127     127  /var/ftp                  /sbin/sh
nobody        65534   65534  /
ted           17301     100  /home/ted                 /usr/bin/csh
+                 0       0
```

The columns correspond to various fields of the **passwd** file. The last entry, **+**, indicates that the system also recognizes login accounts in the NIS database. See "Adding Login Accounts" for more information on **passwd** entry format.

Setting and Listing Login Account Defaults

The Add operation of the Login Account menu supplies default values for queries where appropriate. To list the current defaults, use the operation `User -> Login Account -> Defaults -> Get`.

If these defaults do not suit your needs, you can change them with the operation `User -> Login Account -> Defaults -> Set`.

You can set and list defaults for the following login account parameters.

Group This is the user group to which the new user will belong. A user may belong to multiple groups (changing groups as desired with the **newgrp**(1)

command), but you may supply only one user group name or ID number as the default group. The user group must already exist. Operations for managing user groups are in the Group menu.

Base Directory

This is the directory that will contain the user's home directory. By default, the base directory is **/home**, so a new user's home directory is **/home**/*username*.

Skeleton Directory

This directory contains sample or prototype configuration files, such as **.profile**, **.login**, **.cshrc**, and so on. If you wish, you may create your own skeleton directory. Do not add or change files in the system default skeleton directory because future revisions of the DG/UX system may overwrite it with new contents. The files in the skeleton directory provide a foundation from which the user may begin to customize his or her own operating environment.

Shell Program

The choice of shell program typically depends on personal preference. The three choices are:

sh(1) The Bourne shell.

csh(1) The C shell.

ksh(1) The Korn shell.

For a comparison of the shells, see *Using the DG/UX™ System*. As an alternative, you may specify another initial program.

Managing User Groups

Including users in *groups* allows you to grant them certain file and directory privileges while excluding users outside of the group. Use the **chmod**(1)and **chgrp**(1) commands to control group privileges in the file system.

A user may belong to multiple groups, but a user's shell process is associated with only one group at a time. To change one's current group affiliation, use the **newgrp**(1) command.

Local user group information is in the file **/etc/group**. Each line in the file represents a user group and gives the user group name, the group ID number, and names of members in the group.

The Group menu provides operations for adding, deleting, modifying, and listing user groups.

Adding User Groups

Select the Add operation to add an entry for a new user group. The operation adds the entry to the local **/etc/group** file unless your system is the master server in an NIS domain, in which case it lets you choose whether to add the entry to the local **group** database or to the global NIS database.

A **group** entry consists of several fields, described below:

Group Name
The name may be up to 32 alphanumeric characters long, starting with a letter. To remain compatible with more traditional systems, however, you may want to limit group names to eight characters. The group name must not already exist, either in your local **group** file or, where applicable, in the global NIS database.

Group ID
The group ID, or GID, is a number between 0 and 60,000, inclusive. The numbers below 100 are reserved for system use. The GID must not already be assigned to a group name, either in your local **group** file or, where applicable, in the global NIS database. The default GID is the highest assigned GID plus one, as long as it is less than 60,000; if no such number exists, the default is the lowest unassigned number.

Group Members
List the login names of the users who will be members of the group. Separate them with commas. The login names must already exist as entries in your local **passwd** file or, where applicable, in the global NIS database.

You must create a user group before trying to add login accounts for users who will be members of the group.

Deleting User Groups

Select the Delete operation to remove a user group. The operation removes the group from the local **/etc/group** file unless your system is the master server in an NIS domain, in which case it lets you choose whether you want to remove the group from the local **group** file or from the global NIS database.

After you delete a group, the system can no longer resolve references to the group ID number and produce the group name. This means that commands such as **ls**(1) will list the group ID number instead of the name, where applicable. If you add another group and give it the ID number of an old group, any files associated with the old group will now appear associated with the new one.

Modifying User Groups

Select the Modify operation to change an entry in the **/etc/group** file. The operation changes the local **group** file unless you are the master server in an NIS domain, in which case the operation lets you choose whether to change the local **group** file or the global NIS database.

The operation first prompts you for the name of the group you wish to change. The group must already exist. Then the operation prompts you with:

Group ID
The default is the group's current GID. If you wish to change the number, enter another. The ID number may not already exist. A GID must be between 100 and 60,000, inclusive.

Group Member Modification

For changing the list of group members, this query offers four choices:

no change Select this value if you do not wish to change the membership.

append Select this value to add users to the group. At the next prompt, **Group Members**, specify the login names to append to the current membership list.

remove Select this value to delete users from the group. At the next prompt, **Group Members**, specify the login names to remove from the current membership list.

replace all Select this value if you wish to replace the entire list of members with a new list. At the next prompt, **Group Members**, specify the new list of login names.

At the next prompt, **New Group Name**, you may specify a new name for the group if you wish.

Changes become visible in the local **group** file immediately. If you change the global NIS database, changes may not be visible until the following day.

Displaying User Groups

Select the List operation to display entries from the **/etc/group** file. The List operation gets its information from the local **group** file unless your system is in an NIS domain, in which case you may choose between the local **group** file and the global NIS database.

An example local groups listing follows.

```
Group                  Gid   Members
-----                ------   -------
root                     0   root
other                    1
bin                      2   root, bin, daemon
sys                      3   root, bin, sys, adm
adm                      4   root, adm, daemon
mail                     5   mail, bin
lp                       6   lp
uucp                     8   uucp
daemon                  12   root, daemon
operator                18   adm
nfs                     38   nfs
ftp                     39   ftp
general                100
+                        0
```

The display shows the group name, the group ID number, and the membership list. The last entry, +, indicates that the system also recognizes group accounts in the NIS database. See "Adding Group Accounts" for more information on **group** entry format.

Managing Mail Aliases

A mail alias is a name that stands for one or more login names or mail aliases. The **sendmail**(1C)-based mail facility of the DG/UX system uses aliases when resolving mail addresses.

Mail addresses are useful for two primary purposes:

- A mail alias that resolves to a list of names can act as a mailing list so that any mail sent to the alias name goes out to all names in the alias. For example, you can have an alias called **muleskinners** that resolves to login names **willie**, **eldupree**, **paco**, and **jed**.
- A mail alias that resolves to one name can provide an easy way of redirecting mail sent to a general-purpose address. For example, you can set up the alias **wizard** so it directs mail to the resident computer expert at your facility. If the resident expert gets tired of answering user questions or simply runs out of answers altogether, you can change the **wizard** alias to point to another system expert.

The Mail Alias menu of the User menu provides operations for adding, deleting, modifying, and listing mail aliases. The system keeps track of mail aliases in the **/etc/aliases** file. The master server in an NIS domain also tracks mail aliases.

Adding Mail Aliases

Select the Add operation to add a mail alias to your system. The operation adds the mail alias to your local **/etc/aliases** file unless your system is the master server of an NIS domain, in which case you may choose whether to add the entry to the local **aliases** file or to the global NIS database.

The Add operation first prompts you to supply an alias name. The name can be up to 32 characters long, consisting of upper- and lowercase letters, numerals, and the hyphen (-). The first character should be an alphabetic character. The name must be unique among aliases on your system and, if applicable, in your NIS domain.

The operation also prompts you for an optional list of names. The names should be existing login names or existing alias names. You can include a login name or a mail alias in multiple mail aliases.

The new alias becomes available in the mail system immediately if you are changing the local **aliases** file. If you are changing the global NIS database, the alias may not be available until the following day.

The Add operation uses the **newaliases**(1C) command to make the new alias available.

Deleting Mail Aliases

Select the Delete operation to remove an alias from the alias database. The operation removes the alias from the local **/etc/aliases** file unless your system is the master server of an NIS domain, in which case you may choose whether to remove the alias from the local **aliases** file or from the global NIS database.

Once you have deleted a mail alias, the mail system considers mail addressed to the alias to be misdirected mail. The mail system then returns the mailed message to the sender or forwards the mail to the **root** mail file.

Modifying Mail Aliases

To change the name of an alias or to change the list of names to which an alias resolves, select the Modify operation. The operation changes the local **/etc/aliases** file unless your system is the master server of an NIS domain, in which case the operation lets you choose whether to change the local **aliases** file or the global NIS database.

The operation first prompts you for the name of the alias you wish to change. The alias must already exist. The operation also allows you to list the alias before entering a change.

For changing the list of names in the alias, the next prompt, **Member Modification**, offers four choices:

no change Select this value if you do not wish to change the alias list.

append Select this value to add names to the alias list. At the next prompt, **Alias Member(s)**, specify the names to append to the current alias list.

remove Select this value to delete names from the alias. At the next prompt, **Alias Member(s)**, specify the names to remove from the alias list.

replace all Select this value if you wish to replace the entire alias list with a new list. At the next prompt, **Alias Member(s)**, specify the new list of alias names.

At the prompt, **New Alias Name**, you may specify a new name for the alias if you wish.

Changes become visible in the mail environment immediately. If you change the global NIS database, changes may not be visible until the following day.

Displaying Mail Aliases

Select the List operation to display mail aliases. If your system is in an NIS domain, you may choose whether you want to display aliases from the local **/etc/aliases** file or from the global NIS database.

An example alias listing follows.

```
Mail Alias              Mail Address(es)
----------              ----------------
MAILER-DAEMON           root

postmaster              root

aunts                   polly
                        jemima
                        grizelda

cowpokes                bubba@blarney.state.edu
                        sales03!wilbur@acme.com
                        george@grumble.mil
```

The User's Environment

The profile is the key element in establishing an environment in which users can be the most productive.

Among the things that a profile can contain are:

- A PATH (searchlist) specifying directories to be searched for commands.
- Definitions of TERM (terminal) and TZ (time zone) variables.
- A command that prints the message-of-the-day file, **/etc/motd**, upon login. See Chapter 2 for more information on **motd**.
- Commands to print notification of new mail messages.

Types of Profile

Profiles are of two types: global and local.

The Global Profile

The global profile is an ASCII text file, **/etc/profile** for **sh** and **ksh** users or **/etc/login.csh** for **csh** users, that helps to set up the user's working environment. The global profile contains commands, shell procedures, and environment variable assignments. When a user logs in, the **login** process executes the global profile for the user's initial shell. Users may further customize their personal environment by creating local profiles.

Figure 13–1 shows the global **profile** file shipped with DG/UX.

```
#!/bin/sh
#
# Copyright (C) Data General Corporation, 1984 - 1988
# All Rights Reserved.
# Licensed Material-Property of Data General Corporation.
# This software is made available solely pursuant to the
# terms of a DGC license agreement which governs its use.
#
# Rcsid = "<@(#) profile.proto.sh,v      5.1.1.1>"
#
#       <PassStamp:_@(#)DG/UX_5.4.2__e2.0:5.7>
#
#       The profile that all logins get before using their own
.profile.

trap ""  2 3

#       Set LOGNAME
export LOGNAME

#       Set TZ.
if [ -f /etc/TIMEZONE ]
then
       /etc/TIMEZONE
elif [ -f /etc/TIMEZONE.proto ]
then
       /etc/TIMEZONE.proto
fi

#       Set TERM.
if      [ -z "$TERM" ]
then
        TERM=vt100        # for standard async terminal/console
        export TERM
fi

#       Login and -su shells get /etc/profile services.
#       -rsh is given its environment in its .profile.
case "$0" in
-su )
        export PATH
        ;;
-sh )
        export PATH

        #       Allow the user to break the Message-Of-The-Day only.
       trap "trap " 2"  2
       if [ -f /usr/bin/cat ] ; then
              cat -s /etc/motd
       fi
       trap "" 2

       if [ -f /usr/bin/mail ] ; then
              if mail -e ; then
                     echo "you have mail"
              fi
       fi
       ;;
esac

#       set the umask for more secure operation
#umask 022
trap  2 3
```

Figure 13–1 /etc/profile: Global Profile for sh and ksh Users

Figure 13–2 shows the global **login.csh** file shipped with DG/UX.

```
#
# Copyright (C) Data General Corporation, 1984 - 1988
# All Rights Reserved.
# Licensed Material-Property of Data General Corporation.
# This software is made available solely pursuant to the
# terms of a DGC license agreement which governs its use.
#
# Rcsid = "<@(#) login.csh.proto.csh,v  5.1.1.1>"
#
# /etc/login.csh
#       The .login that all users get before using their own .login.

#       Set LOGNAME
setenv USER   $LOGNAME

#       Set TZ.
if ( -r /etc/TIMEZONE.csh ) then
        source /etc/TIMEZONE.csh
endif

#       If term not set or null, set it to vt100.
if ( ! $?term ) then
    set term
endif
if ("$term" == "" ) then
        set term=vt100
        setenv TERM $term
endif

#       Protect us from ourselves.
set noclobber ignoreeof.
```

Figure 13–2 /etc/login.csh: Global Profile for csh Users

The Local Profile

The local or individual profile is **.profile** for **sh** users, and **.login** for **csh** users. These files reside in a user's home directory. These local profiles are copies of two prototype files: **/etc/skel/.profile** and **/etc/skel/.cshrc**. At the local profile level, users can add commands and variables to customize their environments. A local profile does not have to exist, but if it does exist, it is executed at login time, after the execution of the global profiles **/etc/profile** or **/etc/login.csh**.

Environment Variables

An array of strings called the environment is made available by **exec**(2) when a process begins. Since **login** is a process, the array of environment strings is made available to it. The local profiles in Figure 13–3 show how environment variables can be defined for users running the Bourne shell (**sh**(1)) or the C shell (**csh**(1)).

```
#LOCAL .PROFILE (sh or ksh)    #LOCAL .LOGIN (csh)
#                              #
#                              #
LOGNAME=poulet                 set prompt = 'sys5>'
PATH=:/bin:/usr/bin            set noclobber
MAIL=/usr/mail/poulet          setenv PATH $HOME/bin:/usr/bin:/bin
```

Figure 13–3 Local Profiles

Other programs make use of the information in the environment array list. New strings can be defined at any time. For a discussion of shell syntax, *Using the DG/UX™ System* or the **sh**(1) or **csh**(1) manual pages.

Default Permissions Mode: umask

A system default controls the permissions mode of any files or directories created by a user. The DG/UX system gives default values of 666 for files and 777 for directories (see **chmod**(1M)). The default for files gives all users read and write permission. For directories, all users get, write, and execute permission. Execute permission on a directory lets you make the directory your working directory and list its contents.

If you consider the default permissions too permissive, you can set your own defaults by including the **umask**(1) command, with some appropriate argument, in your personal profile. The **umask** command alters your default permission levels by the amount specified in the argument. The argument is a three-digit number where the first digit tells how much to reduce the digit representing the owner's permissions; the second digit tells how much to reduce the digit representing group permissions; and the third digit tells how much to reduce the digit representing permissions for others.

The following line, for example, would reduce the default owner permissions by 0, the default group permissions by 2, and the default permissions for others by 7:

```
umask 027
```

Thus, the resulting default for directory creation is to set permissions to 750, and for file creation, 640. See the **chmod**(1) manual page for more information on permissions.

Default Shell and Restricted Shell

Generally, when a user logs in the default program that is started is **/sbin/sh**. There may be cases, however, where a user needs to be given a restricted shell, **/bin/restsh**. A restricted shell is one where the user is not allowed to:

- Change directories
- Change the value of PATH
- Specify pathnames or command names containing a slash (/). That is, the user of a restricted shell may not access files or directories other than the present working directory or those included in PATH
- Redirect output

You can use a restricted shell strategy to limit certain users to the execution of a small number of commands or programs. By setting up a special directory for executables and controlling PATH so it only references that directory, you can restrict a user's activity in whatever way is appropriate for your system.

Expert Information

The remainder of this chapter provides additional information on user services. The **sysadm** program should take care of any questions you have while you are performing the major procedures for user services. Still, you should feel free to read and use this additional information to enhance your understanding and performance.

For more information on setting up terminals and keyboards to accommodate your environment, see *Customizing the DG/UX™ System*.

User Passwords

Before users are permitted to log in to your system, they must be listed either in the local **/etc/passwd** file or, if your system is in an NIS domain, in the global NIS **yppasswd** file. An entry in the **passwd** file consists of a single line with the following seven colon-separated fields:

login_name:*password*:*uid*:*gid*:*comment*:*home_directory*:*program*

For a user named L.Q. Poulet, the line might look like the following:

```
poulet:Rm27oQak1:103:104:L.Q. Poulet,,,:/home/poulet:/usr/bin/csh
```

The fields are:

login name: poulet
: A valid name for logging onto the system that is usually chosen by the user. A login name can be up to 32 alphabetic or numeric characters. To remain compatible with traditional systems, you may choose to limit login names to eight characters. Also, by default only the first eight characters of a login name are checked for uniqueness, so any names longer than eight characters must not have the same first eight characters.

password: Rm27oQak1
: A user chooses a password and registers it with the system with the **passwd**(1) command. The encrypted form of the password appears in this field. No one but the user ever knows the real password. The actual password can be a maximum of eight characters. At least one character *must* be a numeric character or special character. This policy discourages users from choosing ordinary words as passwords. When you add a user to the file, you may use a default password, such as **passwd9**, and instruct the user to change it at the first login. Following the encrypted password, separated by a comma, there may be a field that controls password aging. For information on password aging see **passwd**(4) and **passwd**(1).

user ID: 103
: The user ID number (**uid**) is between 100 and 60,000. The number may not include any punctuation. Numbers 99 and below are reserved. User ID 0 is reserved for the superuser.

group ID: 104
: The same conditions apply to the group ID (**gid**) number as to the **uid**.

comment: L.Q.Poulet,,,
: Optional. May contain user's name, office, office phone number, home phone number, and so on. Strictly speaking, there is no required format for this field. For historical reasons, this field is also called the **gecos** (JEE-cose) field. Some utilities, such as **finger**, expect the **gecos** field to be in a particular format. See the **finger**(1) manual page.

home directory: /home/poulet
: The directory where the user is placed upon logging in. The name is usually the same as the login name, preceded by a parent directory such as **/home**. The home directory is the origination point of the user's directory tree.

program: /bin/csh
: The name of a program invoked at the time the user logs in. If the field is empty, the default program is **/sbin/sh**. This field is most commonly used to invoke a special shell, such as **/bin/restsh** (restricted shell).

As noted above, the password field may contain a subfield that controls the aging of passwords. A description of how the process works can be found in **passwd**(4). The effect is to force users periodically to select a new password. If password aging is not used, a person can keep the same password indefinitely.

Password Aging

The password aging mechanism forces users to change their passwords periodically. It also prevents them from changing a password before a specified time interval. You can select password aging for a login account when you execute `User -> Login Account -> Add`.

The password aging information is appended to the encrypted password field in the **passwd** file. The password aging information consists of a comma and up to four characters in the format:

,Mmww

The meaning of these fields is as follows:

, The delimiter between the password itself and the aging information.

M A single character from the 64-character alphabet (described below) representing the maximum duration of the password in weeks.

m A single character from the 64-character alphabet (described below) representing the minimum time interval before the existing password can be changed by the user, in weeks.

ww Two characters from the 64-character alphabet (described below) representing the week (counted from the beginning of 1970) when the password was last changed. You add this information through the codes that you edit into the **passwd** file. Then, the system automatically adds these

characters to the password aging field. All times are specified in weeks (0 through 63) by a 64-character alphabet: **.** (dot), / (slash), **0–9**, **A–Z**, **a–z**; where **\.** is zero, / is one, **0** is 2, and so on. The expression **Ea**, for example, represents the base-10 value 1062.

Table 13–1 shows the relationship between the numerical values and character codes. Any of the character codes may be used in the four fields of the password aging information.

Table 13–1 Password Aging Codes

Character	Number of Weeks
. (period)	0 (zero)
/ (slash)	1
0 through **9**	2 through 11
A through **Z**	12 through 37
a through **z**	38 through 63

Two special cases apply for the character codes:

- If *M* and *m* are equal to zero (the code **.**), the user is forced to change the password at the next login. No further password aging is then applied to that login.
- If *m* is greater than *M* (for example, the code /), only the superuser (**sysadm** or **root**) is able to change the password for that login.

Group IDs

Group IDs are a means of establishing another level of ownership of and access to files and directories. Users with some community of interest can be identified as members of the same group. Any file created by a member of the group carries the group ID as a secondary identification. By manipulating the permissions field of the file, the owner (or someone with the effective user ID of the owner) can grant read, write, or execute privileges to other group members.

Information about groups is kept in the **/etc/group** file. Entries consist of the following colon-separated fields:

```
group_name:password:gid:login_names
```

A sample entry from this file is shown and explained below:

```
prog::104:reynard,poulet
```

Each entry is one line; each line has the following fields:

group_name: prog
The group name can be up to 32 characters, although you may want to limit

it to 8 characters to remain compatible with other systems. The first character must be alphabetic.

password:
The password field should not be used. Leave this field blank.

group_ID: 104
The group ID is a number from 100 to 60,000. The number may not include any punctuation. Numbers below 100 are reserved.

login_names: reynard,poulet
The login names of group members are in a comma-separated list. A user may be a member of more than one group. Nothing prevents a user from having more than one login name, however, as long as each name is unique within the system.

Sample /etc/passwd Entries

Password administration can be set up in a variety of ways to meet the needs of different organizations. The following example shows the password aging information required to establish a new password every 2 weeks (**0**) and to deny changing the new password for 1 week (/).

Here is a typical login/password entry in the **/etc/passwd** file for the typical user **bas**:

```
bas:mst3kmMOE2m.E:100:1:Bo Smith,,,:/home/bas:/usr/bin/csh
```

To require **bas** to change the password at least every 2 weeks, but keep it at least for 1 week, you should use the code **0/**. After you edit the **passwd** file, adding **,0/** to the password field, the entry looks like this:

```
bas:mst3kmMOE2m.E,0/:100:1:Bo Smith,,,:/home/bas:/usr/bin/csh
```

After the password entry is changed, **bas** will have to change the password at the next login and every 2 weeks thereafter.

After **bas**'s first login following the change, the system automatically adds the two-character, "last-time-changed" information to the password field.

```
bas:mst3kmMOE2m.E,0/W9:100:1:Bo Smith,,,:/home/bas:/usr/bin/csh
```

In this example, **bas** changed the password in week **W9**.

Changing or Deleting Aliases

As with changes to the **/etc/passwd** file, all changes that you make when adding or deleting an alias by hand in the **/etc/aliases** file should adhere to a format that the **sysadm** program can use. This format is:

```
alias-name:          name1, name2, name3, name4, name5, name6,
                     name7, name8, name9, name10

alias-name2:         nameA, nameB, nameC, nameD
```

There must be a colon after each `alias-name`. All alias member names, except the last one, must be separated by commas; spaces are ignored. The last entry in the member list should not be followed by a comma. If `name10` had a comma after it, the system would search for another member name and would erroneously read `alias-name2` as a member name. After you edit **/etc/aliases**, run the following command:

```
# /usr/bin/newaliases ↵
```

This command initializes the alias database, displaying the total number of aliases, the length in bytes of the longest alias, and the length in bytes of the entire aliases file. When the command has finished, your changes are available to the system.

End of Chapter

Chapter 14
Accounting

Through **sysadm**, you can print or display on your terminal system-use data that is logged and then stored in summary files and reports. These reports are useful for keeping track of system use by login, command, or machine.

This section summarizes the tasks involved in using the accounting system. The remainder of this chapter shows you in detail how the components of the accounting system function.

Starting and Stopping

By default, accounting is turned off. To start the accounting system, execute the **sysadm** operation `System -> Accounting -> Start`. To turn accounting off, execute `System -> Accounting -> Stop`.

Turning on accounting starts accounting immediately. It also arranges to have accounting run when the system reboots and arranges for the following accounting functions to be run on a regular basis with **cron**:

```
0 4 * * * /bin/su - adm -c "/usr/lib/acct/runacct 2> \
   /var/adm/acct/nite/fd2log"
5 * * * * /bin/su - adm -c "/usr/lib/acct/ckpacct"
15 5 1 * * /bin/su - adm -c /usr/lib/acct/monacct
0 2 * * * /usr/lib/acct/dodisk
```

The default jobs are described below:

runacct	Runs at 04:00 a.m. to collate statistics on connects, user activity, CPU usage, fees, disk usage, and so on. Error messages go to **/var/adm/acct/nite/fd2log**.
ckpacct	Runs at 5 minutes after every hour to monitor the size of the **pacct** data repository file, renaming it when it reaches its size limit (by default, 500 blocks).
monacct	Runs at 05:15 (5:15 in the morning) on the first day of every month to collate and summarize monthly statistics. The **monacct** program cleans up all daily reports and daily total accounting files and deposits one monthly total report and one monthly total accounting file in the **fiscal** directory. The default action of **monacct** adds the current month's date to the file names.
dodisk	Runs at 02:00 a.m. to perform disk accounting.

To change the default values for accounting **cron** jobs, execute `System -> Accounting -> Cron -> Change`, which invokes an editing session with an editor such as **vi**, allowing you to hand-edit the **cron** accounting prototype file.

Listing the Accounting Reports

You can obtain accounting reports with the **sysadm** operation `System -> Accounting -> List`, from which you select one of these report types:

- Command usage
- User logins
- Full report

The full report has the following sections:

- Daily line usage
- Daily usage by login name
- Daily total command summary
- Monthly total command summary
- Last login

Each time you request a full listing, you generate all of these reports automatically. You cannot generate a single report.

Daily Line Usage

The first part of the daily line usage report is the **from/to** banner. The banner displays the time the last accounting report was generated and the time the current accounting report is generated. It is followed by a log of system reboots, shutdowns, recoveries, and any other record dumped into **/etc/wtmp** by the **acctwtmp** program.

The second part of the report is a breakdown of line usage. `Total Duration` tells how long the system was accessible through the terminal lines.

The following is an example of a daily line usage report:

```
Nov 11 16:00 1993  DAILY REPORT For DG/UX page 1

from Tue Nov 10 08:27:00 1993
to   Wed Nov 11 04:00:42 1993

2    shutdown

1    runacct

Total Duration is 1174 Minutes

Line    Minutes    Percent    # Sess    # On    # Off
tty01         0          0         1       1        0

tty02       506         43         1       1        1

tty03        48          4         5       5        5

console      41          3         1       1        1

TOTALS      693         --         9       9        8
```

The columns in the report above are defined as follows:

`Line`	Terminal line or access port
`Minutes`	Total minutes of line use during the accounting period
`Percent`	Total minutes of line use divided into the total duration
`# Sess`	Number of times this port was accessed for a **login**(1) session
`# On`	Number of times the port was used to log a user in to the system.
`# Off`	Number of times a user logged off, and any interrupts on that line.

Generally, interrupts occur on a port when a port service is first enabled on a port when the system goes to multi-user mode. Interrupts occur at a rate of about two per event. Therefore you often see more than twice the number of `# Off` than `# On` or `# Sess`. If the number of `# Off` exceeds the number of `# On` by a large factor, it usually indicates a faulty or failing multiplexer, modem, or cable connection. An unconnected cable dangling from the multiplexer can cause this.

Daily Usage by Login Name

The following is an example report for each user:

```
Nov 11 04:42 1993 Daily Usage Report

Uid Login     CPU (Min)    Kcore-Mins  Connect(Min)  Disk    # Of  # Of # Disk
    Name    Prime NPrime Prime NPrime Prime NPrime  Blocks Procs Sess Samples Fee

0   TOTAL   37    235    1198  467190 56    130     0      6142  11   0       0

0   root    5     1      472   84     41    0       0      712   1    1       0

2   bin     0     0      0     0      0     0       397415 0     0    1       0

3   sys     0     0      0     0      0     0       0      0     1    0       0

4   adm     0     1      0     78     0     0       17630  291   0    1       0

5   uucp    2     2      182   249    0     0       0      740   2    1       0

8   mail    0     0      13    3      0     0       0      13    4    1       0

201 moe     5     0      718   0      0     0       0      151   1    1       0

202 larry   0     1      15    5      0     99      0      13    1    1       0

203 curly   25    5      3069  371    15    31      3345   404   2    1       0

204 poulet  0     0      0     0      0     0       32     0     0    1       0
```

The daily usage by login name report gives a breakdown of system resource use for each user. Its data consists of:

`Uid`	User ID.
`Login Name`	Login name of the user. There can be more than one login name for a single user ID.
`CPU (Min)`	CPU time used, divided into `Prime` and `NPrime` (nonprime). See "Updating holidays" later in this chapter.
`Kcore-Mins`	Total memory a process used, in kilobytes per minute. Divided into `Prime` and `NPrime`.
`Connect(Min)`	How long a user was logged into the system, by `Prime` and `NPrime` time. If this time is long and the column `# of Procs` is low, this user rarely uses the terminal.
`Disk Blocks`	Total amount of disk I/O performed by a user login.
`# Of Procs`	Number of processes invoked by the user. Large numbers may indicate that a shell procedure looped.
`# of Sess`	Number of times the user logged in to the system.

`# Disk Samples`	Number of times the disk accounting was run to obtain the average number of disk blocks.
`Fee`	Total accumulation of total charges against the user. You use the **chargefee**(1M) command to charge a user for special services, such as restoring files and mounting tapes.

Daily and Monthly Total Command Summaries

The daily and monthly total command summaries are similar, but the Daily Command Summary reports only the current accounting period; the Monthly Total Command Summary reports from the start of the fiscal period to the current date. In other words, the monthly summary reflects the data accumulated since the last invocation of **monacct**.

These summaries tell which commands are used most. Since you know which system resources the commands use, you can tune the system accordingly. The summaries are sorted by `Total KCoremin` (see below), which is a good way to calculate drain on a system.

The following are examples of the two types of command summaries. First is an example of a daily command summary:

```
Nov 11 04:42 1993 Daily Command Summary

Command   Numb  Total    Total   Total     Mean    Mean     Hog     Chars   Blocks
Name      Cmds KCoremin CPU-Min Real-Min Size-K  CPU-Min Factor  Trnsfd  Read

TOTALS    2332  1624.74 16.05   15210.91 5.67    003      0.01    0       0

sh        1028  434.53  7.24    7632.44  59.99   0.01     0.01    0       0

csh       1115  474.13  3.96    6534.53  41.111  0.01     0.40    0       0

sendmail  18    55.79   0.93     2.10    155.96  0.04     0.26    0       0

ls        111   62.69   0.57    4.35     98.99   0.01     0.21    0       0

more      35    25.43   0.19    207.16   84.93   0.06     0.00    0       0

ps        1     20.74   0.63    0.35     173.62  0.14     0.33    0       0

cp        20    317.311 2.94    3.41     339.97  0.09     0.21    0       0
```

Next is a monthly command summary:

```
Nov 11 04:42 1993 Monthly Total Summary

Command  Numb   Total     Total   Total     Mean    Mean    Hog    Chars  Blocks
Name     Cmds  KCoremin   CPU-Min Real-Min  Size-K  CPU-Min Factor Trnsfd Read

TOTALS   27792 17118.47  227.94  77021.94  1122.74 3.71    0      0      0

sh       13118 5551.53   93.16   50386.06  59.59   0.01    0.01   0      0

csh      10115 3474.13   33.96   26534.53  41.11   0.01    0.40   0      0

sendmail 238   1455.79   8.93    52.10     165.97  0.03    0.16   0      0

ls       1075  1062.21   9.57    54.85     113.99  0.01    0.17   0      0

more     185   825.43    6.19    107.16    124.93  0.06    0.00   0      0

ps       38    520.74    3.63    11.35     181.72  0.08    0.50   0      0

cp       2970  384.31    8.94    52.85     43.93   0.00    0.17   0      0
```

The columns in the report are defined as follows:

`Command Name`	The name of the command. All shell procedures are reported under the specific shell's name such as **sh**, **csh**, and **ksh**.
`Numb Cmds`	Number of times a command was invoked.
`Total KCoremin`	Total KB segments of memory used by a process, per minute of run time.
`Total CPU-Min`	A program's total processing time.
`Total Real-Min`	Total real-time minutes this program has run.
`Mean Size-K`	Mean of the `Total KCoremin` divided by `Total CPU-Min`.
`Mean CPU-Min`	Mean CPU time (`Total CPU-Min/Numb Cmds`).
`Hog Factor`	Ratio of system availability to system usage (total CPU time/elapsed time). This measures the total available CPU time the process used.
`Chars Trnsfd`	Number of characters manipulated by the **read**(2) and **write**(2) system calls.
`Blocks Read`	Total physical block reads and writes that a process performed.

Last Login

This report gives the date when a particular login name was last used. The report can help you find likely candidates for the tape archives, such as **/usr** directories associated with unused login names.

The following is an example report:

```
Nov 11 04:42 1993 LAST LOGIN

00-00-00        bin     92-11-02        carson

00-00-00        croot2  92-10-09        moe

00-00-00        daemon  92-11-11        larry

00-00-00        svvs    92-11-10        curly

00-00-00        archive 92-11-01        tlp

91-11-12        poulet  92-11-11        uucp
```

A field of all zeros means that login has not been used since the last invocation of the **lastlogin** program.

Updating Holidays

You can update your holiday database with the **sysadm** operation `System -> Accounting -> Edit Holidays File`, which invokes an editing session with an editor, such as **vi**, on the file **/usr/lib/acct/holidays**. The table format has three types of entries:

1. comment lines

Comment lines have an asterisk in column 1.

2. year designation line

This line should be the first data line (a line that is not commented) in the file; it must appear only once. It has three fields: year, prime time, and nonprime time. For example, to specify the year 1992, prime time at 8:30 a.m., and nonprime time at 5:00 p.m., enter:

```
1992  0830  1700
```

The time 2400 is automatically converted to 0000.

3. holiday lines

These entries follow the year designation line and have the format:

```
Day-of-Year  Month Day     Description of Holiday
```

The *Day-of-Year* field is a number between 1 and 366, indicating the day for the corresponding holiday; leading blanks, tabs, and spaces are ignored. The other three fields are commentary and are not currently used by other programs.

The following is an example of a **/usr/lib/acct/holidays** file:

```
*
* Copyright (C) Data General Corporation, 1984 - 1992
* All Rights Reserved.
* Licensed Material-Property of Data General Corporation.
* This software is made available solely pursuant to the
* terms of a DGC license agreement which governs its use.
*
*       $\&What: <@(#) holidays,v 4.1.1.2> $
*
*
* Prime/Nonprime Table for UNIX Accounting System
*
* Curr     Prime     Non-Prime
* Year     Start     Start
*
  1992     0830      1700
*
* Day of   Calendar  Company
* Year     Date      Holiday
*
    1      Jan 1     New Year's Day
  146      May 25    Memorial Day
  185      Jul 3     Independence Day
  251      Sep 7     Labor Day
  331      Nov 26    Thanksgiving
  332      Nov 27    Day After Thanksgiving
  360      Dec 25    Christmas Day
  366      Dec 31    New Years Eve
```

Useful User Accounting Commands

You can invoke the following commands from the shell:

acctcon1 — Reads **/etc/wtmp** and converts login and logoff data to a sequence of records, one per login session. The output is ASCII, giving device, user ID, login name, prime connect time (seconds), nonprime connect time (seconds), session starting time, and starting date and time.

acctdusg — Computes disk resource consumption (including indirect blocks) for each login. See **acct**(1M).

acctwtmp — Records boot times in **/etc/wtmp**.

chargefee — Charges a specified amount against a specified user by generating a charge report and logging it to **/var/adm/fee**. The **runacct** program reads this charge and merges it into the total accounting records.

ckpacct — Checks and controls the size of **/var/adm/pacct**. When **pacct** grows larger than 500 blocks, **turnacct "off"** turns **accton** off, renames the current **pacct** file to **pacct1**, and creates a new **pacct**. Finally, **ckpacct** turns **accton** back on.

dodisk	Does disk accounting on the special files in **/etc/fstab**. Creates **/var/adm/dtmp.**
monacct	Uses the daily data organized by **runacct** and writes it into a monthly summary. Invoke **monacct** through **cron** once each month or each accounting period. **Monacct** creates summary files in **/var/adm/acct/fiscal**. See **acctsh**(1M).
runacct	The main shell program for daily accounting. See **runacct**(1M) for details.
turnacct	An interface to the **accton** program that turns process accounting on or off. When switched on, **turnacct** starts the **accton** program; off stops the **accton** program. When off, the accounting system is disabled.

Recovering from Failure

If the system crashes, **/var** runs out of space, or a **wtmp** file is corrupted, **runacct** fails. If the **active***MMDD* file exists, check it first for error messages. If the **active** file and **lock** file exist, check **fd2log** for error messages. These files are located in **/var/adm/acct/nite**.

Runacct may produce the following error messages. We suggest ways to recover from them.

```
acctg already run for date: check /var/adm/acct/nite/lastdate
```

The date in **lastdate** and today's date are the same. Remove **lastdate**.

```
connect acctg failed: check /var/adm/acct/nite/log
```

The **acctcon1** program encountered a bad **wtmp** file. Use **fwtmp** to fix the bad file. See "Fixing corrupted files" later in this chapter.

```
locks found, run aborted
```

The files **lock** and **lock1** were found in **/var/adm/acct/nite**. Remove these files before restarting **runacct**.

```
Spacct string.MMDD already exists
```

File setups have already run. Check status of files, then run setups manually. See "Restarting runacct" in this chapter for more information.

```
turnacct switch returned rc=string
```

Check the integrity of **/usr/lib/acct/turnacct** and **/usr/lib/acct/accton** by ensuring that the **accton** program is owned by **root** and has the **setuid** bit set. See **chmod**(1).

```
/var/adm/acct/nite/wtmp.MMDD already exists, run setup manually.
```

See "Fixing corrupted files" later in this chapter.

```
wtmpfix errors see /var/adm/acct/nite/wtmperror
```

Wtmpfix detected a corrupted **wtmp** file. Use **fwtmp** to fix the file.

Restarting runacct

Runacct called without arguments assumes this is the first invocation of the day. You must use the argument *MMDD* if **runacct** is being restarted to specify the month and day for which **runacct** will rerun the accounting. The entry point for processing is based on the contents of **/var/adm/acct/nite/statefile**. To override **statefile**, include the desired state on the command line. As mentioned earlier, **runacct** is normally started by **cron**. But should you need to start **runacct** from the command line, here are three ways you might do it.

To start **runacct**, enter:

```
# nohup runacct 2> /var/adm/acct/nite/fd2log &
```

To restart **runacct** specifying *MMDD* (0601), enter:

```
# nohup runacct 0601 2> /var/adm/acct/nite/fd2log &
```

To restart **runacct** at a specific state, such as **wtmpfix**, enter:

```
# nohup runacct 0601 wtmpfix 2> /var/adm/acct/nite/fd2log &
```

In the above examples, the `2>` sends the standard error output to the file named **/var/adm/acct/nite/fd2log**; check this file periodically for error messages. Specifying *MMDD* creates a new **/var/adm/acct/nite/active0601** file dated June 1.

Fixing Corrupted Files

Sometimes, a file is corrupted or lost. Although you may restore some files from backup tapes, you must fix the **/etc/tacct** files yourself.

Fixing wtmp Errors

During normal operation, monitor the size of **/etc/wtmp**, which is the basis for the connect accounting. If the file grows rapidly, execute **acctcon1** to see which tty line is the noisiest. If interruption is occurring at a rapid rate, it can affect general system performance.

Wtmp files record who logged in and when. When the date is changed and the DG/UX system is in multi-user mode, a set of date change records is written into **/etc/wtmp**. The **wtmpfix** program adjusts the time stamps in the **wtmp** records when a date change is encountered.

Some combinations of date changes and reboots, however, result in nonsense lines being added to **/etc/wtmp**; these lines cause **acctcon1** to fail. When this happens, **wtmp***MMDD* is created. The following procedure shows you how to fix the file:

1. Go to directory **/var/adm/acct/nite**.
2. Enter the following command:

```
# fwtmp < wtmpMMDD > xwtmp ↵
```

This command line executes the **fwtmp**(1M) program on the contents of **wtmp***MMDD* and stores the output in file **xwtmp**, effectively converting the binary contents of **wtmp***MMDD* into ASCII format.

3. Use an ASCII editor such as **vi** to delete all remaining binary lines from **xwtmp**.
4. When you're finished editing, convert from ASCII back to binary.

```
# fwtmp -ic < xwtmp > wtmp.MMDD ↵
```

If the **wtmp** file is beyond repair, create a null **wtmp** file by issuing a command like this as superuser:

```
# > /etc/wtmp ↵
```

Cleaning out **wtmp** prevents any charging of connect time. In general, you should check the size of **wtmp** occasionally to see if it is taking up too much space. Reduce the size of **wtmp** either by emptying it as shown above, or by truncating it. To truncate **wtmp**, leaving only the last 3200 characters (the last 50 entries), issue these command lines:

```
# tail -3200c /etc/wtmp > /tmp/wtmp ↵
# mv /tmp/wtmp /etc/wtmp ↵
```

Fixing tacct Errors

If you are using the accounting system to charge users for system resources, the integrity of **/var/adm/acct/sum/tacct** is quite important. Occasionally, corrupt **tacct** records appear with negative numbers, duplicate user IDs, or a user ID of 60,000.

First, check **/var/adm/acct/sum/tacctprev** with **prtacct**. If **prtacct** does not report any errors, patch up **/var/adm/acct/sum/tacct***MMDD* and recreate **sum/tacct**. A simple patch-up procedure is:

1. Go to directory **/var/adm/acct/sum**.
2. Enter the following:

```
# acctmerg -v < tacctMMDD > xtacct ↵
```

This command line executes the **acctmerg**(1M) program on the contents of **tacct***MMDD*, and stores the output in file **xtacct**, effectively converting the binary contents of **tacct***MMDD* into ASCII format.

3. Edit **xtacct** and delete all corrupted records.
4. When you've finished editing, convert from ASCII back to binary. Type the following:

```
# acctmerg -i < xtacct > tacctMMDD ↵
```

Remember that **monacct** removes all the **/var/adm/acct/sum/tacct***MMDD* files. Therefore, when you merge these files together, you recreate **/var/adm/acct/sum/tacct**.

End of Chapter

Chapter 15
Using CD-ROM, Magneto-Optical, Diskette, and Tape Drives

This chapter contains information about the installation and use of the following SCSI drives:

- CD-ROM (Compact Disk-Read-Only Memory) disk drive.
- Multiple-read/write magneto-optical disk drive.
- Diskette drive.
- Tape drive.

In addition, it covers soft errors reported for tape drives.

For additional information on naming and configuring disk and tape drives, see *Customizing the DG/UX™ System*.

General Information

The information in this section applies to all three types of drives: CD-ROM, magneto-optical, and diskette drives.

Make sure you have the proper terminator on the last device in a chain of SCSI devices; otherwise, your system will panic.

If the total length of SCSI cable for a particular adapter is greater than 19 feet, you could get errors while reading or writing to the disks. Excessive cable length could also result in problems when trying to access the last device on the chain.

If you have multiple CD-ROM drives, each must have a unique SCSI ID. Multiple magneto-optical or diskette drives, on the other hand, may be clustered so that as many as four share the same SCSI ID. If you put multiple units on the same SCSI ID, you will need to use device specifications that include a third argument (for unit number). For example, the following specifications represent three 5.25-inch diskette drives at SCSI ID 3:

```
sd(insc(),3,0)
sd(insc(),3,1)
sd(insc(),3,2)
```

For every device on your system, there is a short name created at boot time that you can use to refer to the device. You can find a table showing the long name/short

name pairs in the file **/etc/devlinktab**. For example, a workstation's SCSI tape drive has this specification in DG/UX common device specification format:

```
st(insc(),4)
```

With this name in mind, you can locate the device's entry (which has a very similar long name) in the **devlinktab** file. The example below shows the device's entry (as well as the comment lines from the file that have the table headers).

```
# directory     short     long
#
/dev/rmt        0         st(insc@7(FFF8A000),4,0)
```

From this entry, you know that the short name for **st(insc(),4)** is **/dev/rmt/0**. The device is a tape drive, so it has a no-rewind version too, which is **/dev/rmt/0n**.

You use a device's short name when you need to write to it (such as a tape) or mount it as a file system.

The precise adapter specification, SCSI id, and unit number depend on the kind of hardware you have and how the jumpers are set.

Note that when you remove a device and replace it with another device, the new device does not have the same short name as the old device. For example, if you have a tape drive at **/dev/rmt/0** and replace that drive with a new one, the new tape drive has the short name **/dev/rmt/1**.

Every device you add gets a different short name, because the device linking mechanisms cannot know whether the old name and new name are the same type of device, whether the addition of a device is temporary, and so forth. If you want to use the same short device name for a new device, remove **devlinktab** and run **/usr/sbin/init.d/chk.devlink** to create new links.

Using the CD-ROM Disk Device

The DG/UX system supports the following types of file systems on CD-ROM disks for read only operations:

DG/UX
MS–DOS
High Sierra, ISO 9660, or Rock Ridge (also known as CD–ROM)

It is assumed that you have a CD–ROM device that already contains file systems of read–only data. For example, you may have a CD–ROM device that contains online documentation. To make the device's data accessible, you must configure and register the CD–ROM device and add the device's file systems to the **/etc/fstab** file.

If you install or upgrade a DG/UX system containing a CD–ROM device, it is configured automatically in the system's kernel. If you add a CD–ROM device to an operational system, you may dynamically configure the device using the **sysadm** operation `Device -> Configure`. To permanently configure the device, you must

build a custom kernel, entering the device's name in the system file using the **sysadm** operation `System -> Kernel -> Build`. Refer to Chapter 7 for information on dynamic device configuration and Chapter 4 for information on kernel building.

If you install a DG/UX 5.4R3.00 system containing a CD–ROM device, the device is registered automatically. If you upgrade a DG/UX system that supported a release prior to DG/UX 5.4R3.00, the upgrade procedure will not register automatically a CD–ROM device. You must register CD–ROM devices explicitly in compatibility mode using the **sysadm** operation `Device -> Disk -> Physical -> Register`. See Chapter 7 for information on registering a device.

To mount the CD–ROM device's file systems, use the **sysadm** operation `File -> Local Filesys -> Add`. See Chapter 9 for information on adding a file system.

Using the Magneto-Optical Drive

You may use a magneto-optical drive as a tape by writing to it using **tar**(1) or **cpio**(1), or you may use it to contain file systems. When you use a magneto-optical disk to contain file systems, you treat it just like any normal magnetic disk. To prepare a new magneto-optical disk to contain file systems, follow the instructions in Chapters 7 on soft formatting and registering a physical disk and the instructions in Chapter 9 on creating, adding, and mounting file systems.

Before removing a magneto-optical disk from the drive, first unmount any file systems (if mounted) and deregister the device (if registered).

You cannot register this device unless there is a disk in the device.

You cannot use **sysadm** to create a backup on a magneto-optical disk. To make a backup on a magneto-optical disk, use **dump2**(1M). A backup to magneto-optical disk media is limited to one disk; you cannot make a backup with more than one volume.

Using the Diskette Drive

This section contains information specific to diskette drives.

Assigning Unit Numbers

If you intend to install several diskette drives on the same SCSI ID, you need to observe these rules:

- You may have four drives total, units **0** through **3**.
- If the devices on the SCSI ID are all 5.25-inch diskette drives, they may have any of the four available unit numbers (make sure the hardware is jumpered correctly).

- If any of the devices on the SCSI ID is a 3.5-inch diskette drive, then only units 0 and 1 may be used for 3.5-inch drives, and only units 2 and 3 may be used for 5.25-inch drives. If you have only 3.5-inch drives, they are still restricted to units 0 and 1.

NOTE: For Model 6880 diskette drives, you do not assign multiple drives to the same SCSI ID. Each drive is assigned its own SCSI ID. A unit number is not used, so is expressed as 0, by default.

For example:

- If you have three 5.25-inch diskette drives, they can be units 0 through 2.
- If you decide to add a 3.5-inch drive to the same adapter with the three 5.25-inch drives, however, the 3.5-inch drive may be unit 0 or 1, and two of the 5.25-inch drives may be units 2 and 3. The third 5.25-inch drive must use a different SCSI ID.
- If you have three 3.5-inch drives, you may put two on the same SCSI ID, at units 0 and 1, but the third must have a different SCSI ID.

You need to make sure the SCSI ID, unit number, and adapter card straps are set correctly. The adapter card has the SCSI ID settings and the strap settings, and the drive itself has the unit setting.

If your peripheral housing unit contains any 3.5-inch diskette drives, set the straps on the TEAC adapter card to H, F, LEV, STL, and PAR. If the housing unit contains only 5.25-inch diskette drives, set the straps on the TEAC adapter card to H, F, LEV, STL, and PAR. The H is for 1.44 Mbyte and 1.2 Mbyte formats, and F is for 720 KByte format.

If you have two 5.25-inch diskette drives in the same peripheral housing unit, set the IS strap so the drive does not reset the ready state every time the drive changes density between the two 5.25-inch diskette drives.

If you have two 3.5-inch diskette drives in the same peripheral housing unit, set the HHI strap so that high density input is active high when a high density diskette is accessed and active low when a low density diskette is accessed.

Formatting a Diskette

Formatting a diskette is the same as creating a file system on it. Typically, you create a virtual disk information table (VDIT) on virtual disks on a physical disk before creating file systems on it. Diskettes are much smaller than typical disks, however, and a VDIT takes up too much space. On a diskette it is more efficient not to create a VDIT and to create a single file system directly on the physical disk instead. To create a file system on a diskette, follow these steps:

1. Use the **mkfs**(1M) command to format the diskette. You need to know the pathname of the physical device, for example, **/dev/pdsk/1**. You may have to look in **/etc/devlinktab** to see how device file names map to device

specifications. If your device does not appear in **devlinktab**, you need to verify that the diskette drive is installed correctly and configured in your kernel before rebooting your system. See Chapter 4 for information to configure a device in your kernel.

By default, **mkfs** makes a DG/UX file system. To create an MS-DOS file system instead, specify the **dos** (or **pc**) option and a density value to **mkfs**. For 5.25-inch diskettes, the supported density values are **360kb** and **1220kb**. For 3.50-inch diskettes, the supported density values are **720kb** and **1440kb**. See the **mkfs**(1M) manual page for complete information on **mkfs** options.

You may also obtain pre-formatted diskettes from your Data General representative.

To access an MS-DOS file system, the MS-DOS file manager driver, **dfm()**, must be configured in your kernel. The system adds automatically an entry for **dfm()** when building an autoconfigured kernel.

2. Add the file systems with `File System -> Local Filesys -> Add`.

   Specify the appropriate file system type when prompted, such as **dos** or **dg/ux**. The device you mount is the same one on which you created the file system.

3. Mount the file system with `File System -> Local Filesys -> Mount`.

Changing Diskettes

Before removing a diskette from the drive, you should unmount the file system. If you remove the diskette without unmounting the file system, you may still unmount the file system even though the diskette is not present.

If you put another diskette in the drive without having unmounted the previous one, however, the system will prompt you for the previous diskette. Until you re-insert the previous diskette and unmount its file system, you may not access the drive.

Using a Diskette as a Tape

Instead of using the diskette as a file system, you may use it as a tape. Like a tape, you can write to the diskette using **tar**, **cpio**, **dump**, **dump2**, or **dd**. Refer to the device as **/dev/rpdsk** or **/dev/pdsk**. Before you can use the device this way, you must unmount it (if mounted) and deregister it (if registered).

Because diskettes do not have tape marks (as found on magnetic tape), you cannot use multiple diskettes with **tar**, **dump**, **dump2**, or **dd**. The **cpio** command, on the other hand, allows you to use multiple diskettes when archiving files.

You cannot use **sysadm**, that is, the `File System -> Backup -> Create` operation, to dump to a diskette drive. You may dump to a diskette from the shell using the **dump2**(1M) command, but you are limited to one diskette.

If the diskette drive is registered as a file system when you try to write to it as a file (with **tar** or **cpio**), you will receive the error, `Conflict on open`. Deregister the device before writing to it with **tar** or **cpio**.

Recognizing Soft SCSI Tape Drive Errors

Soft tape errors are reported on the system console for SCSI tape drives. These error reports are deferred, meaning the system does not report them as they occur. Instead, the system compares the number of bytes transferred and the number of errors that occurred, and reports according to the ratio of errors to bytes. If there are a large number of errors, the system reports an acceptance level of marginal. If there are a very large number of errors, the system reports an acceptance level of bad.

The DG/UX system logs these error messages, which by default goes to the system console and to **/usr/adm/messages**.

An example of a marginal message follows:

```
Feb  1 17:00:03 blatz dg/ux: Tape device at st(ncsc@7(FFFB0000,7),3,0)
encountered a high number of correctable (soft) errors.
Feb  1 17:00:03 blatz dg/ux: Please observe the suggested maintenance
schedule for the drive.
```

Suggested maintenance includes cleaning the drive heads.

An example of a bad message follows:

```
Feb  2 17:20:03 blatz dg/ux: Tape device at st(ncsc@7(FFFB0000,7),2,0)
encountered an unacceptably high number of correctable (soft) errors.
Feb  2 17:20:03 blatz dg/ux: Please clean the tape drive and use a known
good tape.
Feb  2 17:20:03 blatz dg/ux: If you receive this message frequently,
contact your DG service representative.
```

Using Read–only Devices in Compatibility Mode

If your read–only devices—CD–ROM and WORM (write once read many) disks—contained logical disks, starting in DG/UX 5.4R3.00, they must now be registered in compatibility mode. Since read–only devices cannot be formatted, logical–disk formatted physical devices cannot be converted to virtual disk format. Its data can be read but not altered with operations that expand, delete or re–create the disk's metadata. Operating in compatibility mode neither affects your access to those files or data nor reduces the performance of those file systems.

If you upgrade a DG/UX system that supported a release prior to DG/UX 5.4R3.00, the upgrade procedure will not register automatically a CD–ROM device. You must register CD–ROM devices explicitly in compatibility mode using the **sysadm** operation `Device -> Disk -> Physical -> Register`. See Chapter 7 for information on registering a device.

To verify the read–only devices that operate in compatibility mode, perform the **sysadm** operation `Device -> Disk -> Physical -> List`. The physical disk listing indicates either a virtual disk or logical disk layout.

Should you choose to convert a read–only device to virtual–disk format, you must first perform a physical disk copy to a writable device using the **sysadm** operation `Device -> Disk -> Physical -> Copy`.

To convert the copied physical disk to virtual disk format, use the **sysadm** operation `Device -> Disk -> Physical -> Convert`.

End of Chapter

Chapter 16
Logging System Errors

This chapter tells you how to maintain a record of system and network errors. **Sysadm** automates a procedure for using the system and network error logs.

Using the System Error Log

With the system error log you establish the conditions under which system errors are collected, identify the logging repositories, and generate error log reports.

Turning on Logging

Select `Logging-> System-> Add` to enable specific types of error logging. The Add operation presents you with these prompts:

`Facility Producing Error`

Enter the origin of the message being logged. The default is **`user`**. Valid values are:

`user`	Messages generated by user processes; the default.
`kern`	Messages generated by the kernel.
`mail`	The mail system.
`daemon`	Daemon system servers such as **ftpd**.
`auth`	The authorization system: **login**, **su**, and **ttymon**.
`lpr`	The printer spooling system.
`cron`	The **cron** or **at** facility.
`local0-7`	Reserved for local use.
`mark`	For timestamp messages produced internally by **syslogd**.
`news`	Reserved for the USENET network news system.
`uucp`	Reserved for the UUCP system.
`all`	Indicates all facilities except the mark facility.

`Error Severity Level`

Enter the severity level. Valid values, in descending order of severity, are:

`emergency`	Panic conditions that normally would be broadcast to all users.
`alert`	Conditions that should be corrected immediately, such as a corrupted system database.

critical Messages about critical conditions, such as hard device errors.

error Other errors.

warning Warning messages.

notice Conditions that are not caused by an error, but may require attention.

information Informational messages.

debug Messages that normally are used when debugging a program.

none No messages.

`Logging Destination:`

Enter where you want to forward the message to:

save to file The system prompts for the file name, which begins with a leading slash (/).

write message to user(s)
The system prompts for the username of the recipient (such as root) or a list of user names separated by commas. The message is written to the screens of all recipients in the list who are currently logged in.

send to remote host
The system prompts for the remote host name.

Deleting Log Selections

Through Delete, you can delete a logging selection by the facility producing the error, severity level, and logging destination for a particular host. Select `Logging -> System -> Delete`. The Delete operation presents you with this prompt:

`Entry to Delete:`
There is no default. Enter the selection by facility, level, and action to be deleted. Type **?** for a list, from which you can select entries for deletion.

Modifying Log Selections

Through Modify, you can change any of the settings previously made when you enabled logging. Select `Logging -> System-> Modify` to alter error logging. The Modify operation presents you with this prompt:

`Entry to Modify:`
Enter the system logging entry to be changed in the database. Type **?** to get a current list, an example of which follows:

```
Choices are

     1  *.err            /dev/console
     2  kern.debug       /dev/console
     3  auth.notice      /dev/console
     4  *.err            /usr/adm/messages
     5  kern.debug       /usr/adm/messages
     6  daemon.notice    /usr/adm/messages
     7  auth.notice      /usr/adm/messages
     8  mail.crit        /usr/adm/messages
     9  kern.crit        /var/adm/dgsvcmgr/log.com
    10  user.info        /dev/console
    11  *.alert          root
    12  daemon.info      /var/adm/daemon.info
    13  *.emerg          *
    14  daemon.notice    /var/adm/log/admd.log
    15  daemon.notice    /var/adm/messages

Enter a number, a name, the initial part of a name,
<NL> to take the default,? for help, ^ to return to
the previous query, < to restart the operation,or q
to quit.
```

Supply the number of the entry to modify. Remaining prompts are identical to those presented for the Add operation.

Listing System Log Selections

Through List, you can display the list of facilities and severity levels that are being logged. Select `Logging -> System -> List`. Sample output from this selection follows:

```
Facility    Level       Destination
------      -----       -----------------------
*           error       file /dev/syscon

kernel      debug       file /dev/syscon

auth        notice      file /dev/syscon

*           error       file /usr/adm/messages

kernel      debug       file /usr/adm/messages

daemon      notice      file /usr/adm/messages
auth        notice      file /usr/adm/messages

mail        critical    file /usr/adm/messages

kernel      critical    file /var/adm/dgsvcmgr/log.com

user        info        remote host jester
*           alert       write to root

daemon      info        file /var/adm/daemon.info

*           emergency   all logged-in users
```

An asterisk (*) selects all facilities.

Generating a Log Report

The Report operation generates a report of system errors that have been logged to files on the target host. Select `Logging -> System -> Report`. The Report operation presents you with these prompts:

`Originating Host(s):`
: Enter the names of the hosts whose messages you want to see. Only local log files are consulted, but those files may include messages sent from other hosts by way of remote host entries. To specify a list, separate each host name with a comma. The default is all hosts.

`Identifier(s):`
: Enter the names of the programs that report the system error. To specify a list, separate each identifier name with a comma. The default is all identifiers.

`Message(s):`

Allows you to provide a regular expression to match the type of error messages you are interested in. You may need to view all messages so that you can derive a meaningful pattern. For example, to see only NetWorker-related errors, you can enter this regular expression:

Net.*

to match this pattern:

```
NetWorker savegroup: info patriot_full (with one client)
```

Sample output follows:

```
Date            Host       Identifier   Message
--------------  ---------  -----------  ------------------
Feb 19 23:49:35  patriot    dg/ux        Tape device at
                                      st(cisc@28(FFFFF300,7),5,0
Feb 19 23:49:35  patriot    dg/ux        encountered a hard
                                       error at block
                                      0,status = 4005007
Feb 19 23:48:15  patriot    syslog       NetWorker media:
                                       (notice) 8mm tape
                                 patriot_week3_a used 274 MB of 2000
Feb 19 23:48:10  patriot    dg/ux        Tape device at
                                     st(cisc@28(FFFFF300,7),5,0)
Feb 19 23:48:10  patriot    dg/ux        encountered a hard
                                       error at block
                                      0, status = 4005007
Feb 19 09:52:37  patriot    syslog       NetWorker index:
                                       (notice) cross-
                                      check has completed.
Feb 19 09:52:23  patriot    syslog       NetWorker index:
                                       (notice) nsrck has
                                      completed.
Feb 19 09:52:12  patriot    syslog       NetWorker Server:
                                       (notice) started
Feb 19 09:51:03  patriot    dg/ux        Firmware in SCSI
                                       controller
                                 cisc@28(FFFFF300,7) is out of date
Feb 19 09:51:03  patriot    dg/ux        see release notice.
Feb 19 09:51:03  patriot    dgsvcd       AV/Alert System:
                                      Disabled
Feb 19 09:47:06  patriot    syslogd      going down on
                                      signal 15
Feb 19 09:47:05  patriot    dgsvcd       AV/Alert System:
                                      Going down on
                                      signal 15
```

Using the Network Error Log

With the network error log you establish the conditions under which network error messages are listed and deleted. The errors reported are dependent on the software that is installed. LAN device drivers, TCP/IP, X.25, or OSI can report such errors.

Select `Logging -> Network -> List` to view the network error logs. The List operation presents you with these prompts:

`Age (in days):`
: Enter the age (in days) of the oldest logged message to view. The default is three days. The default would present messages up to and including three days old, including one and two days old.

`Facility Producing Error:`
: Enter the origin of the message being logged. The list produced depends on the software that is installed. By default, messages from all sources are listed. Type **?** for a list of sources.

`Error Severity Level:`
: Enter the severity level; recognized values, in descending order of severity, are as follows:

emergency	Panic conditions that normally would be broadcast to all users.
alert	Conditions that should be corrected immediately, such as a corrupted system database.
critical	Messages about critical conditions, such as hard device errors.
error	Other errors.
warning	Warning messages.
notice	Conditions that are not error conditions, but may require attention.
information	Informational messages.
debug	Messages that normally are used when debugging a program.

Deleting Log Messages

Through Delete, you can delete log messages by age. Select `Logging -> Network -> Delete`. The Delete operation presents you with a single prompt:

`Age (in days):`
: Enter the age (in days) of the logged messages you want to keep. The default is three days. All messages older than the specified age will be deleted.

End of Chapter

Appendix A
fsck Error Conditions

The **fsck**(1M) program checks the internal consistency of file systems, repairing inconsistencies when invoked to do so. At boot, the system runs **fsck** on all local file systems according to options specified in the **fstab** file. For more information on file systems and how to use **fsck** to check them, see Chapter 9.

When **fsck** detects an inconsistency, it reports the error condition to the operator. If a response is required, **fsck** prints a prompt message and waits for a response. This appendix explains the meaning of each error condition, possible responses, and related error conditions.

In "Error Messages for Phased Checking," the error conditions are organized by the phase of the **fsck** program in which they can occur. The error conditions that may occur in more than one phase are discussed under "General Error Messages."

Error messages that occur only during fast recovery file system checking (when checking a file system that had **fsck** logging turned on) appear "Error Messages Exclusive to Fast Recovery Checking" at the end of the appendix.

The following error messages are presented in their basic form. Fatal errors (such as during **fsck -p**) cause the error message to be prefaced by the string `Fatal Error:`. Running with **-p** also causes messages to be preceded by the name of the file system to which the message applies. The following abbreviations appear in the description of error messages:

B A (decimal) disk block number.

N A decimal number.

O An octal number.

C A character.

D, *F* A directory name, file name or pathname string.

I An inode description string. At the very least, this will consist of the inode number. If possible, the inode's size, file type, file mode, UID, GID, time of last modification, owner name, group name and pathname will also be present.

Error Messages for Phased Checking

The **fsck** program checks file systems in phases only if **fsck** logging was not turned on for the file system. This section lists errors you may see during the typical, phased **fsck** check. Some of these error messages may also appear during checking of a file system for which **fsck** logging was turned on.

The section at the end of this appendix lists errors that can appear only during checking of file systems for which **fsck** logging was turned on.

General Error Messages

The messages described in this section may appear at any time during an **fsck** session.

Cannot allocate memory for internal tables (N bytes requested)

The **fsck** program cannot allocate enough memory; this can only occur during stand-alone **fsck**. The **fsck** program will abort. Bring up your system and use the **fsck** command instead.

Cannot read block B

A disk read of block number *B* has failed. The **fsck** program treats the block it could not read as if it were filled with all zeroes, and continues execution, but the file system is not marked as mountable upon conclusion of checking. Use **diskman** to remap the bad block *B* and run **fsck** again.

Cannot write block B

A disk write of block number B has failed. The **fsck** program continues execution, but the file system being checked is not marked as mountable upon conclusion of checking. Use **diskman** to remap the bad block *B*, and run **fsck** again.

Fork failed

The **fsck** program has failed in an attempt to spawn a child process. This will only occur when running **fsck** with the **-p** option. The only file system affected will be the one for which the child **fsck** process was being created; no check will occur.

Internal Software Error: Cannot seek to block B — aborting

A disk seek to block number *B* has failed; this should never happen. Contact your Data General support representative if this message is displayed. The **fsck** program terminates.

Invalid response; please answer yes or no

An invalid answer has been entered in response to one of **fsck**'s questions. The **fsck** program will not continue until a valid response has been entered. The following strings are valid responses: **y, Y, yes, YES, n, N, no** and **NO**.

Errors During fsck Invocation

Before starting to check a file system, **fsck** must parse its command line and determine which files to check. The following messages result from command line errors or information in the file **/etc/fstab**.

The directory D is the mount point for F

The **fsck** program has been given a directory *D* to check and has determined that *D* is the mount point for the file system *F*. This message is purely advisory.

The flags -y, -n, -p, -q and -S are all mutually exclusive

More than one of the above flags has been specified on the **fsck** command line. At most one of them is allowed. The **fsck** program will abort.

Unknown option: -C

An unknown option flag, *C*, has been specified on the **fsck** command line. Valid flags are as follows: **-l**, **-y**, **-n**, **-p**, **-q**, **-t**, **-D**, **-f**, **-s**, **-S**, and **-x**. When you give it an invalid option, **fsck** will abort.

Errors During fsck Initialization

Before a file system check can be performed, **fsck** must set up certain tables and open certain files. The following messages can result from errors during this phase.

Block B is invalid Inode Table Block — rewrite as empty block?

The inode table block *B* does not contain the proper self-identification information. If run with the **-p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to rewrite the block.

Possible responses to the `rewrite as empty block?` prompt are:

YES Fix this error condition by rewriting this block as an empty file node table block. Any inodes that formerly occupied slots in this block will be cleared.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Cannot determine disk size of F

The **fsck** program has been given a file system *F* to check, but the size of *F* cannot be determined. This should never happen. Contact your Data General support representative if this message is displayed. The **fsck** program will abort checking this file system.

Cannot find a readable copy of the superblock

Neither of the two copies of the superblock can be read. The **fsck** program will abort checking this file system.

Cannot find a valid copy of the superblock

Neither of the two copies of the superblock contain the required self-identification information. The **fsck** program will abort checking this file system.

Cannot open F for reading

The **fsck** program has been given a file system *F* to check, but *F* cannot be opened for reading. Check the mode of *F*. The **fsck** program will abort checking this file system.

Cannot open F for writing

The **fsck** program has been given a file system *F* to check, but *F* cannot be opened for writing. Check the mode of *F* and make sure that no disks containing the file

system are physically write-disabled. The **fsck** program will abort checking this file system.

Cannot read superblock copy N

One of the two superblock copies cannot be read. The **fsck** program will attempt to use the other copy and continue.

F is not a regular file, block–special file, character–special file or valid mount point

The **fsck** program has been given a file system *F* to check, but *F* is not of the correct type. *F* must be a file of type ordinary, block-special or character-special, or else it must be listed in the file **/etc/fstab** as a valid mount point directory. The **fsck** program will abort checking this file system.

File system is too large to check

Stand-alone **fsck** cannot allocate enough memory for its internal tables to begin checking the file system. The **fsck** program will abort checking this file system. Bring up your system and use the **fsck** command instead.

File system size stored in superblock is incorrect (N1 blocks should be N2) — fix?

The superblocks contain an incorrect file system size figure. If run with the **–p** or **–q** options, **fsck** will automatically correct this. Otherwise, **fsck** will ask to correct the size.

Possible responses to the `fix?` prompt are:

YES Fix this error condition by setting the file system size to N2, the actual size of the disk containing the file system.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Invalid default Data Element Size exponent: N — fix?

The default data element size for files (stored in the superblocks as a base-2 logarithm) is invalid. If run with the **–p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to set the size's exponent to the default of 4 (meaning an element size of 16 blocks).

Possible responses to the `fix?` prompt are:

YES Fix this error condition by setting the default data element size's exponent to 4.

NO Ignore this error condition. The **fsck** program will abort checking this file system.

Invalid default Directory Data Element Size exponent: N — fix?

The default data element size for directories (stored in the superblocks as a base-2 logarithm) is invalid. If run with the **–p** option, **fsck** will abort checking this file

system. Otherwise, **fsck** will ask to set the size's exponent to the default of 4 (meaning an element size of 16 blocks).

Possible responses to the `fix?` prompt are:

YES Fix this error condition by setting the default directory data element size's exponent to 4.

NO Ignore this error condition. The **fsck** program will abort checking this file system.

Invalid default Directory Index Element Size exponent: N — fix?

The default index element size for directories (stored in the superblocks as a base-2 logarithm) is invalid. If run with the **-p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to set the size's exponent to the default of 0 (meaning an element size of 1 block).

Possible responses to the `fix?` prompt are:

YES Fix this error condition by setting the default directory index element size's exponent to 0.

NO Ignore this error condition. The **fsck** program will abort checking this file system.

Invalid default Index Element Size exponent: N — fix?

The default index element size for files (stored in the superblocks as a base-2 logarithm) is invalid. If run with the **-p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to set the size's exponent to the default of 0 (meaning an element size of 1 block).

Possible responses to the `fix?` prompt are:

YES Fix this error condition by setting the default index element size's exponent to 0.

NO Ignore this error condition. The **fsck** program will abort checking this file system.

Invalid Disk Allocation Region size: N blocks

The DAR size stored in the superblocks is invalid. The **fsck** program will abort checking this file system.

Invalid first allocation threshold file size: N — fix?

The superblocks contain an invalid first allocation threshold file size (the number of blocks a file can allocate in its initial DAR before all subsequent allocations are made from a different DAR). If run with the **-p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to correct the size.

Possible responses to the `fix?` prompt are:

YES Fix this error condition by setting the first allocation threshold file size to the default limit for DARs of the size specified in the superblock.

NO Ignore this error condition. The **fsck** program will abort checking this file system.

Invalid number of inodes per Disk Allocation Region

The number of inodes per DAR stored in the superblocks is invalid. The **fsck** program will abort checking this file system.

Invalid second allocation threshold file size: N — fix?

The superblocks contain an invalid second allocation threshold file size (the number of blocks a file can allocate in a noninitial DAR before all subsequent allocations are made from a different DAR). If run with the **-p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to correct the size.

Possible responses to the `fix?` prompt are:

YES Fix this error condition by setting the second allocation threshold file size to the default limit for DARs of the size specified in the superblock.

NO Ignore this error condition. The **fsck** program will abort checking this file system.

No check necessary for F

The file system F is already marked mountable and **fsck** was invoked with the **-x** flag. The **fsck** program will not check this file system.

Superblock copies differ; using newer copy

Both copies of the superblock are readable and both contain the required self-identification information, but they differ. The **fsck** program will use the first copy (which is guaranteed to be more recent) and continue.

Superblock copy N is invalid

One of the two superblock copies does not contain the required self-identification information. The **fsck** program will attempt to use the other copy and continue.

Superblock has invalid contents in reserved area — fix?

A copy of the superblock has nonzero contents in a reserved area. If running with the **-p** flag, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to fix the reserved area.

Possible responses to the `fix?` prompt are:

YES The superblock's reserved area is initialized so that it contains all 0s.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Errors During Phase 1 – Check Blocks and File Sizes

This phase of **fsck** operation is concerned with inodes. The following messages result from errors in inode types, inode format, file sizes and the data element pointers and index element pointers that make up a file's structure.

Incorrect block count in Inode I (N1 should be N2) — fix?

The inode *I*'s count of the blocks it uses is incorrect. If run with the **-p** option, **fsck** will automatically correct the count to *N2*. Otherwise, **fsck** will ask to correct the count.

Possible responses to the `fix?` prompt are:

YES Fix this error condition by setting inode *I*'s block count to *N2*.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Inode I claims an invalid block (B) — clear bad pointer?

The inode I claims block B, which does not exist. If run with the **-p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to clear the element pointer claiming the invalid block.

Possible responses to the `clear bad pointer?` prompt are:

YES Fix this error condition by clearing the pointer in inode I that claims the nonexistent block. This may result in a "hole" in the file if the cleared pointer was before the last block of the file.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Inode I claims a system block (B) — clear bad pointer?

The inode I claims block B, which is a system block (a bitmap block, file node table block, DAR entry table block or superblock). If run with the **-p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to clear the element pointer claiming the system block.

Possible responses to the `clear bad pointer?` prompt are:

YES Fix this error condition by clearing the pointer in inode I that claims the system block. This may result in a "hole" in the file if the cleared pointer was before the last block of the file.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Inode I has an Index Block (B) with invalid format — clear bad pointer?

The inode I claims block B as an index block, but block B does not contain the proper self-identification information. If run with the **-p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to clear the element pointer claiming the invalid block.

Possible responses to the `clear bad pointer?` prompt are:

YES Fix this error condition by clearing the pointer in inode I that claims the index block. This may result in a "hole" in the file if the cleared pointer was before the last block of the file.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Inode I has invalid contents in its reserved area — fix?

The inode I does not contain the proper self-identification information. If run with the **–p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to fix the reserved area.

Possible responses to the `fix?` prompt are:

YES Fix this error condition by initializing inode *I*'s reserved area.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Inode I has invalid fragment size exponent (N) — clear?

The inode I has a disallowed exponent representing the size of the file's fragment. If run with the **–p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to clear the file.

Possible responses to the `clear?` prompt are:

YES Fix this error condition by clearing inode I.

NO Ignore this error condition. The **fsck** program will abort checking this file system.

Inode I is of unknown file type (O) — clear?

The inode I is of type *O*, which is an unrecognized octal number. If run with the **–p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to clear the file.

Possible responses to the `clear?` prompt are:

YES Fix this error condition by clearing inode I.

NO Ignore this error condition. The **fsck** program will abort checking this file system.

Inode I is partially truncated — fix?

The inode *I*'s size is shorter than the number of blocks allocated to it. If run with the **–p** option, **fsck** will automatically complete the truncation. Otherwise, **fsck** will ask to complete truncating inode I.

Possible responses to the `fix?` prompt are:

YES Fix this error condition by completing the truncation of inode I down to the size stored in the inode.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Errors During Phase 1b – Resolve Duplicate Claims

When **fsck** finds a block claimed by two or more files, it rescans the file system to find the original claimant of that block. This section lists the error messages that result from settling the claim to the disputed block.

Inode I claims another file's blocks — clear?

The inode I claims some blocks that belong to another file. **fsck** will ask to clear the file.

Possible responses to the `clear?` prompt are:

YES Fix this error condition by clearing inode I.

NO Ignore this error condition. This will result in the same question being asked about the next claimant of the disputed block. As long as enough files are eventually cleared to resolve the duplicate claims on block B, **fsck** will continue normally. However, if at the end of Phase 1b any duplicate claims still exist, **fsck** will not mark this file system as mountable upon completing the check.

Errors During Phase 2 – Check Directory Contents

This phase is concerned with the contents of directories. The messages in this section result from improperly formatted directory blocks, an improperly formatted root directory, and bad directory entries. During this phase, all bad entries and inodes are removed from the file system tree.

Directory inode I has a hole — fix?

The directory inode *I* has at least one "hole" in its file structure (gaps before the end of file). If run with the **-p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to rearrange the directory blocks to fill in the hole.

Possible responses to the `fix?` prompt are:

YES Fix this error condition by rearranging the blocks in the directory to eliminate the hole.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Directory inode I has incorrect child count (N1 should be N2) — fix?

The directory inode *I*'s count of children, *N1*, is incorrect. If run with the **-p** or **-q** options, **fsck** will automatically correct the count to *N2*. Otherwise, **fsck** will ask to correct the child count.

Possible responses to the `fix?` prompt are:

YES Fix this error condition by setting inode *I*'s child count to *N2*.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Directory inode I has an invalid block (B) — rewrite as empty block?

The directory inode *I* has a block (address *B*) which does not contain the proper self-identification information. If run with the **–p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to rewrite the block.

Possible responses to the `rewrite as empty block?` prompt are:

YES Fix this error condition by rewriting block *B* as an empty directory block. Any directory entries that formerly occupied this block will be destroyed.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Directory inode I1 has entry for inode I2 of invalid size — remove bad directory entry?

The directory inode *I1* has a directory entry for inode *I2*, but the entry is too long, too short, or is not a multiple of 4 bytes in size. If run with the **–p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to remove the directory entry for inode I2.

Possible responses to the `remove bad directory entry?` prompt are:

YES Fix this error condition by removing the directory entry for inode *I2*. If inode *I2* is an allocated inode with no remaining links, there will be an opportunity to reattach it in the **/lost+found** directory during Phase 3.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Directory inode I1 has entry for inode I2 which is out of order — remove bad directory entry?

The directory inode *I1* has a directory entry for inode *I2* which has a bad sequence number, meaning that the entry is invalid. If run with the **–p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to remove the directory entry for inode I2.

Possible responses to the `remove bad directory entry?` prompt are:

YES Fix this error condition by removing the directory entry for inode *I2*. If inode *I2* is an allocated inode with no remaining links, there will be an opportunity to reattach it in the **/lost+found** directory during Phase 3.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Directory inode I1 has entry for inode I2 with filename of invalid size — remove bad directory entry?

The directory inode *I1* has a directory entry for inode *I2*, but the entry's file name is too long or too short. If run with the **–p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to remove the directory entry for inode *I2*.

Possible responses to the `remove bad directory entry?` prompt are:

YES Fix this error condition by removing the directory entry for inode *I2*. If inode *I2* is an allocated inode with no remaining links, there will be an opportunity to reattach it in the **lost+found** directory during Phase 3.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Directory inode I1 has entry for inode I2 with an illegal filename: F — remove bad directory entry?

The directory inode *I1* has a directory entry for inode *I2*, but the entry's name *F* is **.** (dot) or **..** (dot-dot). These two names are reserved for the directory's links to itself and to its parent, respectively. If run with the **-p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to remove the directory entry for inode *I2*.

Possible responses to the `remove bad directory entry? prompt are:`

YES Fix this error condition by removing the directory entry for inode *I2*. If inode *I2* is an allocated inode with no remaining links, there will be an opportunity to reattach it in the **lost+found** directory during Phase 3.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Directory inode I1 has entry for inode I2, which has a filename with an illegal character, octal value O — remove bad directory entry?

The directory inode *I1* has a directory entry for inode *I2*, but the entry's name includes the illegal character *O*. A character is disallowed if it is non-ASCII or it is the slash character. If run with the **-p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to remove the directory entry for inode *I2*.

Possible responses to the `remove bad directory entry?` prompt are:

YES Fix this error condition by removing the directory entry for inode *I2*. If inode *I2* is an allocated inode with no remaining links, there will be an opportunity to reattach it in the **lost+found** directory during Phase 3.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Directory inode I1 has entry for inode I2, which has an illegally long pathname — remove bad directory entry?

The directory inode *I1* has a directory entry for inode *I2*, but the pathname for that entry relative to the root of the file system would exceed **MAXPATHLEN** (1024) bytes. If run with the **-p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to remove the directory entry for inode I2.

Possible responses to the `remove bad directory entry?` prompt are:

YES Fix this error condition by removing the directory entry for inode *I2*. If inode *I2* is an allocated inode with no remaining links, there will be an opportunity to reattach it in the **/lost+found** directory during Phase 3.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Directory inode I1 has entry for inode I2, which has invalid contents in its reserved area — fix?

The directory inode *I1* has a directory entry for inode *I2*, which has nonzero information in its reserved area. If run with the **-p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to fix the contents of the reserved area of inode *I2*.

Possible responses to the `fix?` prompt are:

YES Fix this error condition by initializing the reserved area of inode *I2*.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Directory inode I1 has entry for inode number I2, which is invalid — remove bad directory entry?

The directory inode *I1* has a directory entry for inode number *I2*, but *I2* is not a valid inode number. If run with the **-p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to remove the directory entry for inode I2.

Possible responses to the `remove bad directory entry?` prompt are:

YES Fix this error condition by removing the directory entry for inode *I2*.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Directory inode I1 has entry for inode number I2, which is unallocated — remove bad directory entry?

The directory inode *I1* has a directory entry for inode number *I2*, but *I2* is not an allocated inode. If run with the **-p** option, **fsck** will automatically remove the directory entry for inode *I2*. Otherwise, **fsck** will ask to remove the directory entry.

Possible responses to the `remove bad directory entry?` prompt are:

YES Fix this error condition by removing the directory entry for inode *I2*.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Directory inode I1 has entry which is an extraneous link to directory inode I2 — remove bad directory entry?

The directory inode *I1* has a directory entry for inode number *I2*, but *I2* is a directory which does not list *I1* as its parent directory. If run with the **-p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to remove the directory entry for inode I2.

Possible responses to the `remove bad directory entry?` prompt are:

YES Fix this error condition by removing the directory entry for inode *I2*. If inode *I2* is an allocated inode with no remaining links, there will be an opportunity to reattach it in the **/lost+found** directory during Phase 3.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Directory inode I1 has entry which is an extraneous link to symbolic link inode I2 — remove bad directory entry?

The directory inode *I1* has a directory entry for inode number *I2*, but *I2* is a symbolic link which already has another hard link. If run with the **–p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to remove the directory entry for inode *I2*.

Possible responses to the `remove bad directory entry?` prompt are:

YES Fix this error condition by removing the directory entry for inode *I2*.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Directory inode I1 has an entry (for inode I2) which crosses a control point directory boundary — remove bad directory entry?

The directory inode *I1* has a directory entry for inode number *I2*, but *I1* and *I2* have different space parent control point directories. If run with the **–p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to remove the directory entry for inode *I2*.

Possible responses to the `remove bad directory entry?` prompt are:

YES Fix this error condition by removing the directory entry for inode *I2*. If inode *I2* is an allocated inode with no remaining links, there will be an opportunity to reattach it in the **/lost+found** directory during Phase 3.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Incorrect filename length in directory inode I1 for directory inode I2 (N1 should be N2) — fix?

The directory inode *I1* has a directory entry for inode *I2*, but the entry's name length, *N1*, is incorrect. If run with the **–p** option, **fsck** will automatically correct the directory entry's name length to *N2*. Otherwise, **fsck** will ask to correct the name length.

Possible responses to the `fix?` prompt are:

YES Fix this error condition by setting the directory entry's length to *N2* bytes.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Root inode is of wrong file type — fix?

The root inode (inode 2) is not a control point directory. If run with the **-p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to fix the incorrect file type.

Possible responses to the `fix?` prompt are:

YES Fix this error condition by setting the file type of inode 2 to type control point directory.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Root inode is not allocated — fix?

The root inode (inode 2) is not allocated. If run with the **-p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to allocate inode 2.

Possible responses to the `fix?` prompt are:

YES Fix this error condition by allocating inode 2 as the root.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Root inode's parent directory is not the root — fix?

The root inode's parent directory is not the root (itself). If run with the **-p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to list the root inode as its own parent.

Possible responses to the `fix?` prompt are:

YES Fix this error condition by setting the root inode's parent directory to itself.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Root inode's space parent control point directory is not the root — fix?

The root inode's space parent control point directory is not the root (itself). If run with the **-p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to list the root inode as its own space parent.

Possible responses to the `fix?` prompt are:

YES Fix this error condition by setting the root inode's space parent control point directory to itself.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Root inode's space usage limit is too large (N1 should be N2) — fix?

The root inode's space usage limit, *N1*, is bigger than the size of the file system, *N2*. If run with the **-p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to reset the limit to *N2* blocks.

Possible responses to the `fix?` prompt are:

YES Fix this error condition by setting the root inode's space usage limit to *N2* blocks.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Errors During Phase 3 – Check Connectivity

Phase 3 of **fsck** deals with the reconnection of unreferenced files and directories onto the file system tree. The messages in this section result from attempts to connect unreferenced files into the **lost+found** directory. Note also that any of the Phase 2 messages may be seen in this phase, as the contents of any reconnected directories must be checked.

Cannot find enough contiguous free blocks to expand directory inode I

The **fsck** program could not find enough contiguous free blocks to expand the directory inode I. Some unreferenced files may not reconnected as a result of this failure; they can be reconnected during a later **fsck** session after enough space has been freed in the file system.

Control point directory inode I has an entry named 'lost+found' which is not a directory — remove bad directory entry?

The control point directory inode I already has an entry named **/lost+found**, but which is not of type directory. If run with the **-p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to remove the entry.

Possible responses to the `remove bad directory entry?` prompt are:

YES Fix this error condition by removing the bad entry from inode I. The bad entry's inode will itself be reattached in the new **/lost+found** directory which will be created in directory I1.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Could not reconnect inode I

The **fsck** program was unable to reconnect the unreferenced inode I because it could not allocate enough blocks to expand the **/lost+found** directory, or because it could not allocate a free inode to use as the **/lost+found** directory.

Directory inode I is already as large as it can become

The **fsck** program has discovered that a directory it was attempting to expand is already the maximum size a directory can become.

Inode I1 lists as its space parent inode number I2, which is not a valid control point directory — reset space parent to root?

The inode I1 has the noncontrol point directory inode I2 listed as its space parent. If run with the **-p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to reset I1's space parent to inode 2, the root of the file system.

Possible responses to the `reset?` prompt are:

YES Fix this error condition by resetting I1's space parent to inode 2, the root of the file system.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Directory inode I needs to be expanded — fix?

The directory inode I needs to be expanded so that another directory entry can be added to it; I is either the root directory or the **/lost+found** directory. If run with the **–p** or **–q** options, **fsck** will automatically attempt to expand the directory. Otherwise, **fsck** will ask to expand it.

Possible responses to the `fix?` prompt are:

YES Fix this error condition by attempting to expand inode I.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Inode I is unreferenced — clear?

The inode I has no links in the file system and an earlier reconnection failed or was refused.

Possible responses to the `clear?` prompt are:

YES Fix this error condition by clearing inode I. The contents of the file will be destroyed.

NO Ignore this error condition. Inode I will remain unattached and can be reattached during a later **fsck** session provided that enough blocks and/or inodes are free.

Inode I is unreferenced — reconnect?

The inode I has no links in the file system. If run with the **–p** option, **fsck** will automatically attempt to reconnect the file. Otherwise, **fsck** will ask to reconnect it.

Possible responses to the `reconnect?` prompt are:

YES Fix this error condition by reconnecting inode *I* in the **/lost+found** directory, with the name "#*N*", where *N* is the inode number of I.

NO Ignore this error condition.

The lost+found directory inode I already has an entry named 'F' — remove bad directory entry?

The **/lost+found** directory inode I has discovered that it already has an entry of the name F when it was trying to reconnect an unreferenced file which would have had the same name. If run with the **–p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to remove the spurious entry.

Possible responses to the `remove bad directory entry?` prompt are:

YES Fix this error condition by removing the entry for F; the inode referred to by that entry will be reattached with a name constructed from its inode number.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Errors During Phase 4 – Check Link Counts and Resource Accounting

This phase checks the link counts of individual inodes and the resource counts (blocks and inodes used) of control point directories. The messages result from errors in these counts.

Control point directory inode I has incorrect inode allocation count (N1 should be N2) — fix?

The control point directory inode I has a bad count of the inodes used by it and all its space descendants. If run with the **–p** or **–q** options, **fsck** will automatically adjust the count to *N2*. Otherwise, **fsck** will ask to fix the count.

Possible responses to the `fix?` prompt are:

YES Fix this error condition by adjusting the inode count for inode *I* to *N2*.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Control point directory inode I has incorrect space allocation count (N1 should be N2) — fix?

The control point directory inode I has a bad count of the blocks used by it and all its space descendants. If run with the **–p** or **–q** options, **fsck** will automatically adjust the count to *N2*. Otherwise, **fsck** will ask to fix the count.

Possible responses to the `fix?` prompt are:

YES Fix this error condition by adjusting the space count for inode I to *N2*.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Inode I has incorrect link count (N1 should be N2) — fix?

The inode I has a bad link count. If run with the **–p** or **–q** options, **fsck** will automatically adjust the count to *N2*. Otherwise, **fsck** will ask to fix the count.

Possible responses to the `fix?` prompt are:

YES Fix this error condition by adjusting the link count for inode I to *N2*.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Errors During Phase 5 – Check Disk Allocation Region Information

This phase deals with the disk allocation regions. Messages in this section result from errors in the components of the DARs: the bitmap, the free inode list, and various resource counts.

Block B of the Disk Allocation Region Information Area is invalid — fix?

The disk allocation region information area block B does not contain the proper self-identification information. If run with the **-p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to rewrite the block.

Possible responses to the `fix?` prompt are:

YES Fix this error condition by rewriting this block as an empty disk allocation region information area block. The DAR information in the block will be corrected later in this Phase.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Disk Allocation Region N has incorrect Bitmap — fix?

The bitmap for DAR N is incorrect. If run with the **-p** option, **fsck** will automatically correct the bitmap. Otherwise, **fsck** will ask to correct it.

Possible responses to the `fix?` prompt are:

YES Fix this error condition by rewriting the bitmap correctly.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Disk Allocation Region N has incorrect count of blocks used (N1 should be N2) — fix?

The block count for DAR *N* is incorrect. If run with the **-p** or **-q** options, **fsck** will automatically correct the count to *N2*. Otherwise, **fsck** will ask to correct it.

Possible responses to the `fix?` prompt are:

YES Fix this error condition by changing DAR *N*'s block count from *N1* to *N2*.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Disk Allocation Region N has incorrect counts of directories and inodes used — fix?

The counts of used files and directories for DAR N are incorrect. If run with the **-p** or **-q** options, **fsck** will automatically correct the counts. Otherwise, **fsck** will ask to correct them.

Possible responses to the `fix?` prompt are:

YES Fix this error condition by rewriting the counts of used inodes and directories correctly.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Disk Allocation Region N has incorrect free inode list — fix?

The linked list of free inodes in DAR number N is incorrect: it contains allocated inodes, duplicates, or it does not contain some inodes which are actually unallocated. If run with the **–p** or **–q** options, **fsck** will automatically correct the free list. Otherwise, **fsck** will ask to correct it.

Possible responses to the `fix?` prompt are:

YES Fix this error condition by rewriting the free list for DAR number N.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Disk Allocation Region N has invalid contents in its reserved area — fix?

Disk allocation region number N has nonzero contents in its reserved area. If run with the **–p** option, **fsck** will abort checking this file system. Otherwise, **fsck** will ask to zero out the reserved area.

Possible responses to the `fix?` prompt are:

YES Fix this error condition by initializing the contents of the reserved area.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Incorrect summary counts in superblocks — fix?

The counts of used blocks and files in the two copies of the superblock are incorrect. If run with the **–p** or **–q** options, **fsck** will automatically correct the counts. Otherwise, **fsck** will ask to correct them.

Possible responses to the `fix?` prompt are:

YES Fix this error condition by rewriting the counts of used blocks and files correctly.

NO Ignore this error condition. The **fsck** program will not mark this file system as mountable upon completing the check.

Advisory Messages During File System Cleanup

Once a file system has been checked, a few cleanup functions are performed. This section lists advisory messages about the file system.

File System is now mountable

The **fsck** program has successfully completed checking the file system and it has been marked as mountable.

File System is still inconsistent and not mountable

The **fsck** program has completed checking the file system, but inconsistencies remain and the file system is still marked as unmountable. Re-run **fsck** in order to fix the remaining inconsistencies.

N1 of N2 blocks used (N3 free); N4 of N5 inodes used (N6 free)

The indicated number of blocks and inodes have been used, leaving the indicated number unallocated.

Unconnected files still remain. Mount the file system and remove files to free data blocks and inodes

The **fsck** program has successfully completed checking the file system and it has been marked as mountable. However, there are still unreferenced files in the file system. These unreferenced files can be recovered by running **fsck** again after enough blocks and inodes have been freed to allow them room to be reconnected.

Error Messages Exclusive to Fast Recovery Checking

This section lists errors that can only appear during checking of a fast recovery file system, one for which **fsck** logging was turned on. These messages do not appear with phased **fsck** messages in the previous section because fast recovery checking does not occur in phases.

While normal phased checking will return only the messages in the previous section, fast recovery file system checking may return the messages in this section as well as many (but not all) of those appearing in the previous section.

Bitmap block N in Disk Allocation Region N has an invalid format

This message is printed when **fsck** discovers a bitmap block with a bad self-id. The arguments are the lda of the bitmap block and the number of the DAR it belongs to.

Block N is invalid Inode Table block

This message is printed when an Inode Table block fails to self-ID. The numeric argument is the invalid block.

Block N of the Disk Allocation Region Information Area is invalid

This message is printed when **fsck** discovers a DARIA which does not self-ID. The argument is the ordinal number (within the DARIA) of the faulty block.

Cannot allocate memory for internal tables (N bytes requested)

This message is printed when an attempt to allocate memory fails. The numeric argument is the number of bytes requested.

Cannot find enough contiguous free blocks to allocate a missing index element for inode I

This message is printed when **fsck** fails to allocate an index element. The argument describes the inode in question.

Cannot open fast recovery dump file F

This message is printed when the open of the file to hold the human-readable version of the fast recovery log fails. The argument is the pathname of the dump file.

Cannot read DAR table

This message is printed when the DAR table cannot be read.

Cannot read fast recovery log at lda B

This message is printed when the read of one of the fast recovery disk logs fails. The argument is the logical disk address of the log.

Cannot recover from log. Run full fsck

This message is printed when a fatal error happens while trying to recover with the log.

Cannot write DAR table

This message is printed when the DAR table cannot be written.

Directory entry D in Directory inode I contains the wrong inode number (B should be B)

This message is printed when a directory entry is discovered to contain an incorrect file node number. The arguments are the directory entry file name, the file node number of the directory containing the entry, the entry's incorrect file node number, and the correct file node number.

Directory inode I has a spurious entry for file name F with file node number B

This message is printed when an unneeded directory entry is found. The arguments are the directory's file node number, the name of the missing entry and the file node number that should be in that entry.

Directory inode I has an incorrect parent inode number (B should be B)

This message is printed when a directory file node containing the wrong parent inode number in its din is found. The arguments are the directory's inode number and the incorrect and correct parent inode numbers.

Directory inode I has incorrect child count (N should be N)

This message is printed when **fsck** discovers an erroneous child count in a directory. The string argument describes the inode in question. The first numeric argument is the old, incorrect count; the second numeric argument is the correct count.

Directory inode I is missing an entry for file name F with file node number B

This message is printed when a directory entry is found to be missing. The arguments are the directory's file node number, the name of the missing entry and the file node number that should be in that entry.

Disk Allocation Region N has an incorrect Free Inode List

This message is printed when **fsck** discovers a DAR with incorrect Free File Node List. The first argument is the DAR number.

Disk Allocation Region N has incorrect Bitmap

This message is printed when **fsck** discovers an incorrect bitmap. The argument is the number of the DAR with the bad bitmap.

Disk Allocation Region N has incorrect count of blocks used (N should be N)

This message is printed when **fsck** discovers a DAR with incorrect space count. The first argument is the DAR number. The second argument is the incorrect, old count. The third argument is the correct count.

Disk Allocation Region N has incorrect counts of directories and inodes used

This message is printed when **fsck** discovers a DAR with incorrect directory or file node counts. The first argument is the DAR number.

File node and space accounting partially recovered. Run full fsck to recover accounting information

This message is printed when **fsck** completes but accounting recovery from the log was not complete. This error does not prevent the file system from being mounted, but at some point, full **fsck** should be run to fully correct accounting information.

Inode I has a bad fragment size for data element #0 (B should be B)

This message is printed when a file node's `fragment_size_exponent` field is found to contain an incorrect value. The arguments are the inode number, the bad fragment size exponent and the correct fragment size exponent.

Inode I has an Index Block (N) with invalid format

This message is printed when **fsck** discovers an element pointing to an index block which does not self-ID. The string argument describes the inode in question. The numeric argument is the bad block's number.

Inode I has an incorrect file size (N bytes should be N bytes)

This message is printed when **fsck** discovers a node which has an incorrect file size. The string argument describes the inode in question. The first numeric argument is the old, incorrect size. The second numeric argument is the correct size.

Inode I has an incorrect space parent inode number (B should be B)

This message is printed when a file node containing the wrong space parent inode number is found. The arguments are the file's inode number and the incorrect and correct space parent inode numbers.

Inode I has bad element pointers

This message is printed when an element pointer in the file node is found to contain the incorrect address and the element is allocated. The argument is the inode number.

Inode I has incorrect inode allocation count (N should be N)

This message is printed when **fsck** discovers a node which has an incorrect node count. The string argument describes the inode in question. The first numeric argument is the old, incorrect count. The second numeric argument is the correct count.

Inode I has incorrect link count (N should be N)

This message is printed when **fsck** discovers a node which has an incorrect link count. The string argument describes the inode in question. The first numeric argument is the old, incorrect count. The second numeric argument is the correct count.

Inode I has incorrect space allocation count (N should be N)

This message is printed when **fsck** discovers a node which has an incorrect space count. The string argument describes the inode in question. The first numeric argument is the old, incorrect count. The second numeric argument is the correct count.

Inode I is incorrectly marked as allocated

This message is printed when a file node allocation has to be backed out or the results of a deallocation did not make it to disk. The argument is the file node number.

Inode I is partially truncated

This message is printed when **fsck** discovers a partially truncated file. The string argument is the file in question.

Internal software bug O

This message is printed when a sanity check fails, indicating a software bug. The argument is a status encoded with FSCK_ENCODE_BUG_STATUS.

Missing log entry for file system fnode accounting

This message is printed when it is discovered that there are file node allocations or frees on the file system, but that no log entry giving the initial value of the file system's `number_of_used_file_nodes` was made. This error indicates that there is a bug in the kernel logging code.

System buffers containing data from the following files may not have been written to disk:

This message is printed when **fsck** fails to allocate an index element. The argument describes the inode in question.

Unexpected child count N in inode I

This message indicates a bug in the log recovery code.

Unknown kind of log entry

This message is printed when an unrecognized type of log entry is found in the log.

End of Appendix

Appendix B
SAF Reference Cards

This appendix contains tables that you may copy or remove and use as reference cards for the Service Access Facility (SAF). For information on SAF, see Chapter 10.

Port Monitor Management Reference Card

Command Syntax	Description
sacadm -a -p *pmtag* **-t** *type* **-c** *"cmd"* **-v** *ver* \ [**-f dx**] [**-n** *count*] [**-y** *"comment"*] \ [**-z** *script*]	Add a port monitor entry to SAC's administrative file.
sacadm -l [**-p** *pmtag* | **-t** *type]*	Print port monitor status information.
sacadm -L [**-p** *pmtag* | **-t** *type*]	Print port monitor status information in condensed format.
sacadm -G [**-z** *script*]	Print or replace per-system configuration script **/etc/saf/_sysconfig**.
sacadm -g -p *pmtag* [**-z** *script*]	Print or replace per-port monitor configuration script **/etc/saf/**pmtag**/_config**.
sacadm -e -p *pmtag*	Enable port monitor pmtag.
sacadm -d -p *pmtag*	Disable port monitor pmtag.
sacadm -s -p *pmtag*	Start port monitor pmtag.
sacadm -k -p *pmtag*	Kill port monitor pmtag.
sacadm -r -p *pmtag*	Remove the entry for port monitor pmtag from the SAC administrative file.

Service Administration Reference Card

Command Syntax	Description
pmadm -a [**-p** *pmtag* | **-t** *type*] \ **-s** *svctag* **-i** *id* **-m** *"pmspecific"* \ **-v** *ver* [**-f ux**] [**-y** "comment"] \ [**-z** *script*]	Add a service entry to the port monitor administrative file.
pmadm -l [**-t** *type* | **-p** *pmtag*] \ [**-s** *svctag*]	Print service status information.
pmadm -L [**-t** *type* | **-p** *pmtag*] \ [**-s** *svctag*]	Print service status information in condensed format.
pmadm -g -p *pmtag* **-s** *svctag* \ [**-z** *script*]	Print, install, or replace per–service configuration script for service *svctag* associated with port monitor *pmtag*.
pmadm -g -s *svctag* **-t** *type* **-z** script	Install, or replace per-service configuration scripts for all services *svctag* associated with port monitors of type *type*.
pmadm -e -p *pmtag* **-s** *svctag*	Enable service *svctag* associated with port monitor *pmtag*.
pmadm -d -p *pmtag* **-s** *svctag*	Disable service *svctag* associated with port monitor *pmtag*.
pmadm -r -p *pmtag* **-s** *svctag*	Remove the entry for service *svctag* from the port monitor administrative file.

ttymon and Terminal Line Setting Reference Card

Command Syntax	Description
sacadm -l [**-t** *type* | **-p** *pmtag*]	Lists all port monitors (**-l** alone), all port monitors of a given type (**-t** *type*), or a single port monitor (**-p** *pmtag*).
pmadm -l [**-t** *type* | **-p** *pmtag*] \ [**-s** *svctag*]	Lists all services for all port monitors (**-l** alone), all services for all port monitors of a given type (**-t** *type*), all services for a specific port monitor (**-p** *pmtag*), or a single service (**-s** *svctag*).
sacadm -a -p *pmtag* **-t ttymon** \ **-c** *cmd* **-v 'ttyadm -V'**	Adds a ttymon port monitor. **ttyadm** used with **sacadm -a** or **pmadm -a** as an argument to the **-v** option provides the comment line containing the **ttymon** version number for the new port monitor administrative file.
sacadm -r -p *pmtag*	Removes a port monitor.
pmadm -a -p *pmtag* **-s** *svctag* \ **-i** *id* [**-f ux**] **-v 'ttyadm -V'** \ **-m "'ttyadm** [**-b**] [**-r** *count*] \ [**-c**] [**-h**] [**-i** *msg*] [**-m** *modules*] \ [**-p** *prompt*] [**-t** *time-out*] \ **-d** *device* **-l** *ttylabel* \ **-s** *service*'"	Adds a service. **ttyadm** used with **pmadm -a** as an argument to the **-m** option provides the *pmspecific* fields for inclusion in the port monitor's administrative file.
pmadm -r -p *pmtag* **-s** *svctag*	Removes a service.
pmadm -e -p *pmtag* **-s** *svctag*	Enables a service.
pmadm -d -p *pmtag* **-s** *svctag*	Disables the service *svctag*, available through port monitor *pmtag*.
sacadm -e -p *pmtag*	Enables all services defined for port monitor *pmtag*.
sacadm -d -p *pmtag*	Disables all services defined for port monitor *pmtag*.
/usr/sbin/sttydefs -a *ttylabel* \ [**-b**] [**-n** *nextlabel*] \ [**-i** *initial-flags*] \ [**-f** *final-flags*]	Adds an entry to the **/etc/ttydefs** file.
/usr/sbin/sttydefs -l [*ttylabel*]	Prints terminal line setting information from the **/etc/ttydefs** file for terminal ports with the label *ttylabel*. If no *ttylabel* is specified, prints terminal line setting information for all records in the file.
/usr/sbin/sttydefs -r *ttylabel*	Removes records for the *ttylabel* specified from **/etc/ttydefs**.

listen Reference Card

Command Syntax	Description
sacadm -l [**-t** *type*]	Lists status information for all port monitors (**-l** alone) or for all port monitors of a given type (**-t** *type*).
pmadm -l -p *net_spec* [**-s** *svctag*]	If *svctag* is supplied, lists status information for the service. If no service is specified, lists status information for all services under *net_spec*.
sacadm -a -p *net_spec* | *pmtag* \ **-t listen** \ **-c "/usr/lib/saf/listen** *net_spec*" \ **-v 'nlsadmin -V'**	Adds a **listen** port monitor.
sacadm -r -p *net_spec*	Removes a **listen** port monitor.
pmadm -a -p *net_spec* | *pmtag* \ **-s** *svctag* **-i** *id* \ **-m "'nlsadmin** *options*'" \ **-v 'nlsadmin -V' -y** *comment*	Adds a service under a **listen** port monitor.
pmadm -r -p *net_spec* | *pmtag* \ **-s** *svctag*	Removes a service under a **listen** port monitor.
pmadm -e -p *net_spec* **-s** *svctag*	Enables service *svctag* under port monitor *net_spec*.
pmadm -d -p *net_spec* **-s** *svctag*	Disables service *svctag* under port monitor *net_spec*.
sacadm -d -p *net_spec*	Disables all services under port monitor *net_spec*.
sacadm -e -p *net_spec*	Enables all services under *net_spec*.

End of Appendix

Appendix C
DG/UX Directories and Files

This appendix lists the directories and files of the DG/UX system that are of interest to a system administrator. For a list of all directories and files shipped with the DG/UX system, look in the directory **/usr/release/dgux_*.fl**. For detailed information on any of the entries here, see the relevant manual page.

NOTE: You should not modify any of the DG/UX system supplied files or directories. Also, you should not change any of the system supplied directories to symbolic links. Modifying these files and directories may prevent you from upgrading your system to a new revision of DG/UX.

First, this appendix discusses the **root**, **var**, **usr**, and **srv** directories. Then it discusses some files of interest to a system administrator. Notice that some files and directories in the DG/UX system have *physical* locations and *logical* locations. A physical location is a file's real location. A logical location is a symbolic link. For instance, when you reference **/usr/spool/lp** (logical), you are actually referencing **/var/spool/lp** (physical). In this appendix, we denote a symbolic link by enclosing the name in parentheses.

Contents of the Root Directory

The **root** logical disk is mounted on the directory /. We refer to those files on the **root** logical disk as "being in root." There are physical and logical directories on the **root** logical disk. They are:

/.profile

Environment initialization and configuration file for **root** login shell. We recommend that you use the **sysadm** login rather than the **root** login for administrative work.

/admin

This directory is the home directory for the **sysadm** login account. We recommend that you use the **sysadm** login account for administrative work rather than the **root** login so that your work files will be in **/admin** rather than /.

By default, the **/admin** directory contains a prototype shell initialization file, **.profile.proto**, and a working copy of the file, **.profile**. The directory also contains the **crontabs** directory.

The **/admin/crontabs** directory contains prototype **crontab** files useful to the system administrator. The **root.proto** file contains miscellaneous **cron** jobs for maintaining the system. These jobs perform a variety of functions

such as cleaning out temporary file directories and truncating logs. See Chapter 2 for more information. The **lp.proto** file contains jobs for maintaining the LP subsystem. The **uucp.proto** files perform functions for the UUCP file transfer/remote command execution utility. Chapter 2 discusses the **cron** facility and the prototype **crontab** files in more detail.

(/bin) Symbolic link to **usr/bin**. Contains public commands.

/dev Contains device nodes, also called special files.

/dgux Executable kernel file.

/dgux.aviion
Default name of kernel executable file (typically a hard link to **dgux**, above).

/dgux.installer
Executable kernel used for installation of system software.

/etc Contains configuration files and system data.

(/lib) Symbolic link to **usr/lib**. Contains object libraries.

/local Contains site-specific files.

/opt Parent directory for user-configurable portions of applications packages.

/proc Currently unused.

/sbin Contains minimum system commands to get the system up.

/srv A mount point for the **srv** logical disk. Contains client and release management files. See "Contents of the /srv Directory."

/tftpboot
Used only on OS servers and X terminal servers, contains links to OS and X terminal clients' bootable executable files in **/usr/stand**.

/tmp Used for temporary system files.

/usr Mount point directory for the **/usr** file system. See "Contents of the /usr Directory."

/var Used for system data files whose size varies as the system runs. See "Contents of the /var Directory."

Contents of the /var Directory

The **/var** directory contains files that are release dependent, have read and write permissions set, and are dynamically sized.

/var/adm This directory contains a variety of files produced by the system, such as system error logger (**syslogd**(1M)) messages This directory contains the data collection files for the accounting system. See Chapter 14 for complete information on accounting system files and directories.

/var/Build Kernel builds by **sysadm** Auto Configure and Build operations take place here.

/var/cron Contains the **cron** log.

/var/ftp The home directory for **ftp** users. Contains utilities for **ftp** users.

/var/lp Contains logs for the LP print service.

/var/mail Contains **mail** databases and dynamically-sized files.

/var/news Contains **news** databases and dynamically-sized files.

/var/opt Application package parent directory for dynamically-sized files.

/var/saf Contains logs for the Service Access Facility.

/var/spool Contains spooling files for LP, UUCP, and **cron**.

/var/spool/lp Contains all of the files and directories of the LP system. See Chapter 12 for complete information on these files and directories.

/var/spool/cron
Contains **cron** databases and dynamically-sized files.

/var/spool/uucp
Contains files specific to UUCP.

/var/preserve Text editor file save area for sudden program halt.

/var/tmp User temporary file space.

/var/ups Contains logs and the status file for the Uninterruptible Power Supply **upsd**(1M) daemon.

Contents of the /usr Directory

The **usr** logical disk is mounted on the **/usr** directory. Files in this directory are release dependent and read-only. There are physical and logical directories on the **usr** logical disk. They are:

(/usr/adm) Symbolic link to **/var/adm**.

/usr/admin Contains the files, directories, tables, menus, and defaults used by the **sysadm** system administration program.

/usr/bin Contains user commands.

/usr/catman Contains the on-line manual reference pages that users access with the **man**(1) command.

/usr/etc Contains database and configuration files.

/usr/etc/master.d
Contains master files. These files list devices and kernel parameters for the DG/UX system. This directory may also contain master files for other packages that have kernel components, such as TCP/IP and ONC/NFS.

/usr/include Contains include files for system software.

/usr/lib Contains library routines.

/usr/lib/acct Contains the C language programs and shell procedures that drive the accounting system. See Chapter 14.

(/usr/lib/gcc) Symbolic link to **/usr/lib/gcc-2**.

/usr/lib/gcc-2 Contains the DG/UX GNU C compiler. For information see the Release Notice that comes with the compiler package.

/usr/local Contains site-specific, read-only files.

(/usr/mail) Symbolic link to **/var/mail**.

(/usr/news) Symbolic link to **/var/news**.

/usr/opt Contains applications packages.

(/usr/preserve)
Symbolic link to **/var/preserve**.

/usr/release Contains media notices, release notices, and system package names.

/usr/root.proto
Prototype / (root) file system copied when you add an OS client.

/usr/sbin Commands used only by an administrator.

/usr/sbin/init.d
The **init.d** directory contains executable files used to change the system run level. These files are linked to files beginning with **S** (start) or **K** (stop) in **/etc/rc*N*.d**, where *N* is the appropriate run level. Files are only executed from **/etc/rc*N*.d** directories.

/usr/share Contains release independent shared software.

(/usr/spool) Symbolic link to **/var/spool**.

/usr/src Parent directory for source code.

/usr/src/uts/aviion/lb
Contains the kernel libraries which are used to build the kernel image. See Chapter 4. Also see **config**(1M) and **make**(1).

/usr/stand Contains stand-alone utilities and bootstrap programs.

(/usr/tmp) Symbolic link to **/var/tmp**.

(/usr/ucb) Symbolic link to **/usr/bin**.

Contents of the /srv Directory

The **/srv** directory contains the directories and files needed for managing operating system releases and clients.

/srv/admin Contains **sysadm** databases and information files.

/srv/admin/clients
Contains **sysadm** client data.

/srv/admin/defaults
Contains **sysadm** defaults for releases and clients.

/srv/admin/releases
Contains **sysadm** OS release data.

/srv/dump Dump space on a one-per-client basis.

/tftpboot Contains links to bootstraps for diskless clients.

/srv/release Contains space for each release's **usr** and client roots.

/srv/release/PRIMARY
Contains symbolic links to the server's **usr** and / files.

/srv/share Contains release independent shared software.

/srv/swap Swap space on a one-per-client basis.

DG/UX Administrative Files

This section is not intended to be an exhaustive look at the DG/UX files; rather, it highlights the files that you'll be using more often than others. Subsections here are:

- Contents of the **/etc** Directory
- Administrative Commands in the **/sbin** Directory
- Administrative Files in the **/usr** Directory
- Administrative Files in the **/var** Directory

Contents of the /etc Directory

The **/etc** directory contains files that you and your system management tools alter to customize your system. Technically, you may alter most if not all of these files yourself using a text editor; nevertheless, we recommend that in cases where **sysadm** or another administrative utility offers an interface for managing the file, you choose the interface rather than the editor. The rationale for this recommendation is twofold:

- An interface program such as **sysadm** will not corrupt the file it maintains, whereas there is some risk that you may accidentally corrupt the data in the file if you alter it with a text editor.
- An interface program will continue to provide the desired service in future revisions of the software, even if such revisions involve changes such as moving the data file, changing its format, or re-implementing the related service in a completely different manner.

Many of the files in **/etc** and other system directories have *prototypes*, files representing the state of the file as it was shipped with the software release. For example, the prototype for the **passwd** file is **passwd.proto**. You may find the

prototype files useful if you wish to restore or refer to the default configuration of the system. For a complete list of prototype files shipped with the DG/UX system in the **/etc** directory, see **/etc/dgux.prototab**.

/etc Database Files Maintained by the System

This section describes some of the files created and/or maintained by system services other than the **sysadm** utility.

/etc/devlinktab

This files contains entries used to make short-named links to device nodes with otherwise unwieldy names. The system adds to this file at boot to reflect the current configuration of the system. Note that the system does not rebuild this file. Entries are never deleted from **/etc/devlinktab**. An example follows:

```
/dev/rmt        0       st(insc@7(FFF8A000),4,0)
```

The entry above indicates that the tape device with SCSI ID 4 on the integrated SCSI bus appears as device **/dev/rmt/0** in the file system.

/etc/dgux.rclinktab

This data table determines how the **rc.links** script creates or removes the links in the **rc*.d** directories that point to files in the **/usr/sbin/init.d** directory. The DG/UX system and other software packages modify this file during package setup.

/etc/log

A directory containing logs for various system services, including information on run level changes and daemon activity. The system generally creates these logs during the boot process. These logs are useful for reviewing activity that occurs during the boot process and changes of run level.

/etc/rc*.d

There are nine of these directories in **/etc**: **rc0.d**, **rc1.d**, **rc2.d**, **rc3.d**, **rc4.d**, **rc5.d**, **rc6.d**, **rcS.d**, and **rci.d**. Each directory contains links to all the shell scripts in **/etc/init.d**. Software packages modify these directories during setup, reflecting any such changes in **/etc/dgux.rclinktab**. See Chapter 3 for a complete list of all services that can be set for each run level.

/etc/utmp

The **/etc/utmp** file contains information on the run level of the system. Various system services, such as **login**, maintain this information. Use the **who**(1) command to access this information.

/etc/wtmp

The **/etc/wtmp** filecontains a history of system logins. The owner and group of this file must be **adm**, and the access permissions must be 664. This file contains a record for each login that occurs. Use the **last**(1) and **who**(1) commands to access this file. Periodically, this file should be cleared or truncated. See Chapter 2 for information on truncating this file.

/etc Database Files Maintained via sysadm

The **/etc** directory contains a number of files that are databases of information needed to support various subsystems. The following sections describe the files in **/etc** that **sysadm** maintains.

/etc/dgux.params

This file contains parameters that you can set to control the actions of **rc** scripts in **/etc/init.d**. The **rc** scripts determine what happens during boots and other changes of run level. Chapter 3 discusses run levels.

/etc/dumptab

This file contains the dump table which lists the different media supported by **dump2**(1M). It describes the media characteristics for each medium made available to **dump2**. See **dumptab**(4) and Chapter 9 for more information on system backups.

/etc/failover

This directory contains the databases for failover disks. You maintain these databases using the operations in the **sysadm** menu `Device -> Disk -> Failover`. Chapter 9 discusses failover disks.

/etc/fstab

The **fstab** file specifies the file systems to be mounted by the **/etc/mount** command. The following is a sample entry in **fstab**. Note that it is in NFS format; we recommend this format even if you are not using NFS.

```
/dev/dsk/root  /  dg/ux  rw  d  1
```

The above entry indicates a local file system mount, and the following entry indicates an NFS remote file system mount.

```
titan:/usr/titan  nfs  rw,hard  x  0
```

The **fstab** format was changed to support NFS file systems as well as local file systems. The old style **fstab** entries are also supported. See **fstab**(4) and Chapter 9 for detailed information.

/etc/group

The **/etc/group** file describes each group to the system. An entry is added for each new group. Each entry in the file is one line and consists of four fields, which are separated by a colon (:):

group_name:*password*:*group_id*:*login_names*

See "Adding Groups" in Chapter 13 and **group**(4) for more information. If you have ONC/NIS, see **yppasswdd**(1M) and your ONC/NIS documentation.

/etc/inetd.conf

Contains the Internet server configuration database. This is a list of servers that **inetd** invokes when it receives an Internet request over a socket.

/etc/lp This directory contains files supporting your configuration of the LP subsystem. For more information on LP, see Chapter 12.

/etc/nfs.params

This file contains parameters for controlling ONC/NIS and NFS services.

/etc/passwd

The **/etc/passwd** file identifies each user to the system. Add an entry for each new user. Each entry in the file is one line and consists of seven fields. The fields are separated by colons (**:**). The fields are:

login_name:*password*:*uid*:*gid*:*comment*:*home_directory*:*program*

Example:

```
poulet:Rm27oQak1:103:104:L.Q. Poulet:/usr/poulet:/bin/csh
```

See "The User's Environment" in Chapter 13 and **passwd**(4).

/etc/tcpip.params

This file contains the parameters for commands invoked by the **rc** scripts to initialize the network. Chapter 3 describes the **rc** scripts in detail. Also see your TCP/IP documentation.

/etc/TIMEZONE

The **/etc/TIMEZONE** file sets the time zone shell variable TZ. The TZ variable in the **TIMEZONE** file is changed by the **sysadm** operation System –> Date –> Set. The TZ variable can be redefined on a user (login) basis by setting the variable in the associated **.profile** or **.login**. See Chapter 13.

/etc Files Maintained Manually

The **/etc** directory contains a number of files that are databases of information needed to support various subsystems. The following sections describe the files in **/etc** that **sysadm** does not maintain.

/etc/cron.d

This directory contains files that you modify to customize your system's **cron** services. The files here include **at.allow** and **cron.allow**, which list users who are permitted to submit jobs for execution by **cron**(1M), **at**(1), and **batch**(1). The **queuedefs** file, described in more detail in **cron**(1M), contains parameters governing the **at** and **batch** queues. See Chapter 2 and the **crontab**(1) manual page, in addition to the manual pages mentioned above, for more information.

/etc/inittab

The **/etc/inittab** file contains instructions for the **/etc/init** command. The instructions define the processes that are to be created or terminated for each initialization state. Initialization states are called run levels. By convention, run level S is single-user mode; run level 1 is administrative mode; run level 2 is multiuser mode; and run level 3 multiuser mode with network services. Some applications packages (such as X11) modifying this file during package setup. Chapter 3 summarizes the various run levels and describes their uses. See **inittab**(4) for more information.

/etc/login.csh

The default profile for **csh** users is the **/etc/login.csh** file. The default profile for **sh** users is the **/etc/profile** file. The standard (default) environment is established by the commands in these global profile files. See "The User's Environment" in Chapter 13 for more details.

/etc/motd

The **/etc/motd** file contains the message of the day. The system shell initialization files, **login.csh** and **profile**, print this message to users logging in.

/etc/profile

The default profile for **sh** users is in the **/etc/profile** file. The default for **csh** users is the **/etc/login.csh** file. The standard (default) environment is established by the commands in these global profile files. See Chapter 13 for examples.

Administrative Commands in the /sbin Directory

The following commands are available in **/sbin**. These are the minimum system administration commands necessary to operate the system.

/sbin/chk.fsck

An **rc** check script that invokes **fsck**(1M) during boot.

/sbin/fsck

Runs the **fsck**(1M) file system check program.

/sbin/halt

Halts the operating system and restores control to the firmware-based hardware control system, the SCM.

/sbin/init

Command to change run levels S, 1, 2, 3.

/sbin/mount

Mounts a file system on the DG/UX directory tree.

/sbin/rc.init

Executes the shell scripts in **/etc/init.d** via links in **/etc/init/rc*N*.d**. Execution is initiated from entries in **/etc/inittab**. For example, the following line specifies that all scripts associated with run level 3 be executed:

```
rc3:3:wait:/sbin/rc.init 3
```

/sbin/setup.d/boot

Sets up scripts that must run on the host system CPU.

/sbin/setup.d/root

Contains scripts that set up software packages in the / file system.

/sbin/sh

The DG/UX **sh**(1) command.

/sbin/shutdown

Brings the operating system down to single-user mode (run level S). See Chapter 3.

/sbin/su

Switches user (login) name. For instance, **su sysadm** changes your user ID to **sysadm.**

/sbin/ttymon
Enables login on the console line and other terminal lines.

/sbin/umount
Unmounts a remote file system.

Administrative Files in the /usr Directory

The following files are available in the **/usr** directory.

/usr/sbin/init.d
This directory contains the system's check scripts and **rc** scripts.

/usr/sbin/setup.d/usr
This directory contains set up scripts that modify a host's **usr** space. Setup scripts might include those for TCP/IP and ONC/NFS.

/usr/src/uts/aviion/cf/system.*.proto
This file contains your custom version of the devices and configuration parameters listed in **/usr/etc/master.d/dgux**. For configuration, the **config**(1M) program runs on the **system** file and produces program code in **conf.c**. Prototype system files are shipped with software packages with kernel content. Prototype files have names of the form **system.*.proto**.

Administrative Files in the /var Directory

The following files are located in the **/var** directory. These files are useful for monitoring superuser activity and performing recurring system administrative tasks.

/var/adm/sulog
The **/var/adm/sulog** file contains a history of superuser (**su**(1)) command usage. As a security measure, this file should not be readable by others. The **/var/adm/sulog** file should be periodically truncated to keep the size of the file within a reasonable limit. For more information on truncating **sulog**, see Chapter 2.

A typical **/var/adm/sulog** file follows:

```
SU 08/18 16:16 + console smitht-root
SU 08/18 23:45 + tty00 jones-root
SU 08/19 11:53 + console smitht-sysadm
SU 08/19 15:25 + console root-sysadm
SU 08/19 23:45 + tty00 root-uucp
```

/var/cron/log
A history of all actions taken by **/usr/sbin/cron** is recorded in the **/var/cron/log** file. The **/var/cron/log** file should be periodically truncated to keep the size of the file within a reasonable limit. For information on truncating the **cron** log, see Chapter 2.

/var/adm/messages*
This file contains information about the system as logged by **syslog.d**. See Chapter 16 for more information about **syslogd**. The files **messages.0**,

messages.1, and **messages.2** are created by a **cron** job, which is supplied in **/admin/crontabs/root.proto**.

/var/adm/acct
This directory contains the data necessary to provide accounting reports. See Chapter 14 for more information on accounting.

/var/adm/pact
This file is used by the accounting software.

/var/adm/log
This directory contains various log files, which are described in Chapter 2.

End of Appendix

Appendix D
SCSI Disk and Tape Drive Names

This section outlines SCSI disk and tape drive and controller names, including disk-array storage system drive and controller names. It gives only the short form of the device names. For more detail on device names, see *Customizing the DG/UX™ System*.

The format for SCSI disk and tape drives follows.

device (*controller–type* [*@device–code*] , [([*vme*(*n*)] *controller–number* [*,controller–SCSI–ID*), *device–SCSI–ID* [*,LUN*])

Because the format is so elaborate, we show two examples next. The fields in the device name format are explained below. Here are the examples.

Example 1: SCSI disk on ncsc controller

```
                sd(ncsc(0,7),0,0)
                |  |     |  |  |  |
                |  |     |  |  |  |
   device_____  |  |     |  |  |  |___logical-unit-number (LUN)
                   |     |  |  |
controller-type___ |     |  |  |____SCSI-ID-number of disk drive
                         |  |
  controller-number___   |  |___controller SCSI-ID number
```

This device name identifies a disk drive having the SCSI ID 0 on an **ncsc** controller whose SCSI-ID is set to 7. Since the disk is at LUN 0, it needn't be further qualified by a logical unit number, although we do show the 0 for clarity. You could specify this device as **sd(ncsc(),0)**, **sd(ncsc())**, **sd(ncsc(0)**, or **sd(ncsc(0,7),0)**. The first controller number, the first SCSI ID number, and the logical unit number are 0. The SCSI ID of the **ncsc** controller is 7 by default.

Example 2: SCSI disk in disk-array storage system on dgsc controller

```
                sd(dgsc(0),0,2)
                |  |     |  |  |__ disk drive logical unit number
                |  |     |  |     (LUN)(hex)
                |  |     |  |
      device __ |  |     |  |__ SCSI-ID-number of SP
                   |     |     (storage-control processor)
                   |     |
controller-type___ |     |___SCSI-2 adapter controller-number
```

This device name identifies a disk drive in a CLARiiON disk array storage system. The SCSI-2 adapter controller is the first one in the host (number 0), the SP is SP A,

the first one in the storage system (SCSI-ID 0), and the disk drive is logical unit (LUN) 2.

The device name format fields have the following meanings.

device is as a two-letter mnemonic that specifies the disk-array or the standard integrated SCSI device. The valid values, alphabetically, are

- **da** **hada**-type, 30-module disk-array subsystem (this is not the same as a CLARiiON™ disk-array storage system)
- **sd** SCSI disk (on a CSS, PHU, or CLARiiON disk-array storage system)
- **st** SCSI tape (on a CSS, PHU, or CLARiiON tape-array storage system)

controller-type specifies the controller type. Valid values are

- **cisc** Ciprico VME SCSI adapter.
- **dgsc** Data General VME SCSI-2 adapter. When used with a disk-array storage system, it is called a SCSI-2 adapter.
- **hada** H.A.D.A., 30 disk-module disk-array I/O processor.
- **insc** Integrated SCSI adaper.
- **ncsc** NCR integrated SCSI-2 adapter.

device-code is used only in the long form (explained in *Customizing the DG/UX™ System*.)

vme(n) specifies the VME controller channel, for example **(vme(1))**; needed only if the computer has more than one VME channel and this device is on a channel other than the first one.

controller-number indicates the number of this type of controller in the host; controller numbers start at 0 and continue upward.

controller-SCSI-ID is an optional argument. The ID is required for a disk-array SCSI-2 adapter or any other controller in a dual-initiator (shared bus) configuration (described later in this chapter). Table D–1 lists default SCSI IDs.

device-SCSI-ID For a **sd** disk not part of a disk-array storage system or for a **st** tape, this is the SCSI ID number on the disk or tape drive which is jumpered at the factory. Table D–1 lists the default SCSI IDs.

For a disk-array storage system, this field indicates the SP (storage-control processor) SCSI ID. The disk-array storage system can have up to two SPs. The SP SCSI IDs are user selectable, as explained in the 014-series manual that accompanies the storage system.

For a **hada**-type, 30 disk-module disk-array subsystem (**da** device), this field indicates the unit number attached to a physical disk in the disk-array subsystem. A single controller (I/O processor) supports up to 30 disk modules in such a subsystem, numbered 6-23 (hexadecimal).

Table D–1 Default SCSI ID Numbers

Device Type	SCSI ID Number
First disk drive	0
Second disk drive	1
Third disk drive	2
Fourth disk drive	3
First SP (disk-array storage-control processor)	0
First cartridge tape	4
Second cartridge tape	5
First controller or SCSI-2 adapter channel	7
Second controller or SCSI-2 adapter channel	6

LUN For the **sd** and **st** drives, this is the logical unit number of the drive that shares the same SCSI ID. For example, any SCSI ID supports up to 8 unit numbers. A single SCSI ID can manage multiple drives through logical unit numbers, 0-7. The *LUN* field serves primarily for distinguishing disk-array storage system disk drives, diskette drives (except Model 6880 drives; see the Model 6880 manual), and optical disk drives. For the first disk-array storage system on a **dgsc**, you can assign physical unit numbers in the range 0 through 1F hex.

Device Name Example 1

This device name **sd(insc(0),0,0)** identifies a disk drive on the first integrated SCSI controller (insc) which has controller number 0. The disk drive is on SCSI ID 0. Since the disk is at LUN 0, the LUN is not required, although we show it for clarity. You could identify the disk by the name **sd(insc( ),0)** or **sd(insc( ))**. Both the first controller number, the first SCSI ID number, and the logical unit number are 0.

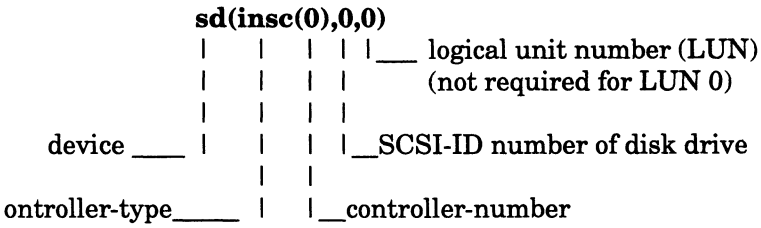

Device Name Example 2

This device name identifies the first diskette drive (LUN 0) attached to a SCSI disk controller (at SCSI ID 3) on the integrated SCSI controller **(insc(0))**.

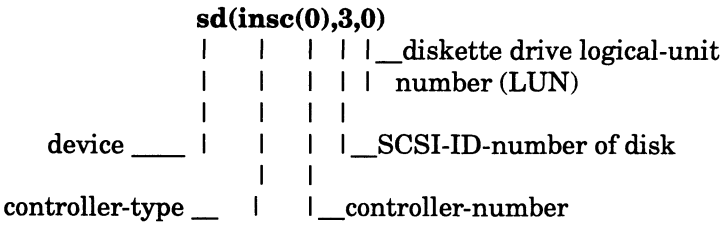

Device Name Example 3

This device name identifies a disk drive in a CLARiiON disk array storage system on the second VME channel. The SCSI-2 adapter controller is the first one in the host's second VME channel [**(dgsc(vme(1),0)**], the SP is SP A, the first one in the storage system (SCSI-ID 0), and the disk drive is logical unit (LUN) number 2.

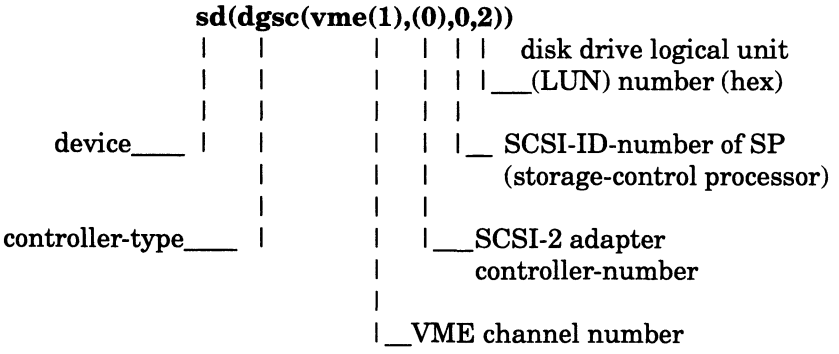

Device Name Example 4

This device name identifies a H.A.D.A. (30-module) disk array, disk unit 12, controlled by the second (1) I/O processor.

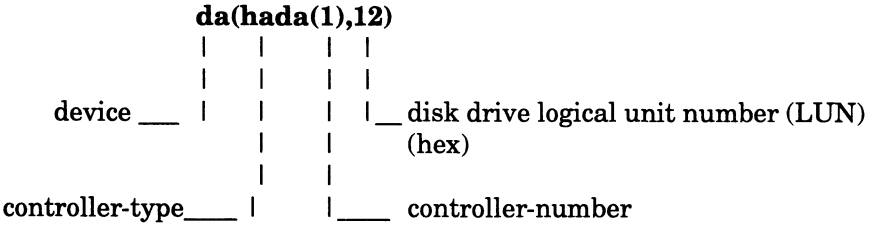

For more detail on device names, see *Customizing the DG/UX™ System*.

End of Appendix

Index

Symbols

Numbers

A

B

C

D

E

F

G

H

I

J

K

L

M

N

O

P

R

S

W

T

U

V

TIPS ORDERING PROCEDURES

TO ORDER

1. An order can be placed with the TIPS group in two ways:
   A. MAIL ORDER – Use the order form on the opposite page and fill in all requested information. Be sure to include shipping charges and local sales tax. If applicable, write in your tax exempt number in the space provided on the order form.

   B. Send your order form with payment to:

   Data General Corporation
   ATTN: Educational Services/TIPS G155
   4400 Computer Drive
   Westboro, MA 01581–9973

   C. TELEPHONE – Call TIPS at (508) 870–1600 for all orders that will be charged by credit card or paid for by purchase orders over $50.00. Operators are available from 8:30 AM to 5:00 PM EST.

METHOD OF PAYMENT

2. As a customer, you have several payment options:
   A. Purchase Order – Minimum of $50. If ordering by mail, a hard copy of the purchase order must accompany order.

   B. Check or Money Order – Make payable to Data General Corporation. Credit Card – A minimum order of $20 is required for MasterCard or Visa orders.

SHIPPING

3. To determine the charge for UPS shipping and handling, check the total quantity of units in your order and refer to the following chart:

Total Quantity	Shipping & Handling Charge
1–4 Items	$5.00
5–10 Items	$8.00
11–40 Items	$10.00
41–200 Items	$30.00
Over 200 Items	$100.00

If overnight or second day shipment is desired, this information should be indicated on the order form. A separate charge will be determined at time of shipment and added to your bill.

VOLUME DISCOUNTS

4. The TIPS discount schedule is based upon the total value of the order.

Order Amount	Discount
$0–$149.99	0%
$150–$499.99	10%
Over $500	20%

TERMS AND CONDITIONS

5. Read the TIPS terms and conditions on the reverse side of the order form carefully. These must be adhered to at all times.

DELIVERY

6. Allow at least two weeks for delivery.

RETURNS

7. Items ordered through the TIPS catalog may not be returned for credit.
8. Order discrepancies must be reported within 15 days of shipment date. Contact your TIPS Administrator at (508) 870–1600 to notify the TIPS department of any problems.

INTERNATIONAL ORDERS

9. Customers outside of the United States must obtain documentation from their local Data General Subsidiary or Representative. Any TIPS orders received by Data General U.S. Headquarters will be forwarded to the appropriate DG Subsidiary or Representative for processing.

TIPS ORDER FORM

Mail To: Data General Corporation
Attn: Educational Services/TIPS G155
4400 Computer Drive
Westboro, MA 01581 - 9973

BILL TO:	**SHIP TO:** (No P.O. Boxes - Complete Only If Different Address)
COMPANY NAME _____	COMPANY NAME _____
ATTN: _____	ATTN: _____
ADDRESS _____	ADDRESS (NO PO BOXES) _____
CITY _____	CITY _____
STATE _____ ZIP _____	STATE _____ ZIP _____

Priority Code _____ (See label on back of catalog)

Authorized Signature of Buyer (Agrees to terms & conditions on reverse side)	Title	Date	Phone (Area Code)	Ext.

ORDER #	QTY	DESCRIPTION	UNIT PRICE	TOTAL PRICE

A SHIPPING & HANDLING

☐ UPS	ADD
1–4 Items	$5.00
5–10 Items	$8.00
11–40 Items	$10.00
41–200 Items	$30.00
200+ Items	$100.00

Check for faster delivery

Additional charge to be determined at time of shipment and added to your bill.

☐ UPS Blue Label (2 day shipping)
☐ Red Label (overnight shipping)

B VOLUME DISCOUNTS

Order Amount	Save
$0–$149.99	0%
$150–$499.99	10%
Over $500.00	20%

Tax Exempt # or Sales Tax (if applicable) _____

ORDER TOTAL	
Less Discount See B	–
SUB TOTAL	
Your local* sales tax	+
Shipping and handling – See A	+
TOTAL – See C	

C PAYMENT METHOD

☐ Purchase Order Attached ($50 minimum)
P.O. number is _____. (Include hardcopy P.O.)
☐ Check or Money Order Enclosed
☐ Visa ☐ MasterCard ($20 minimum on credit cards)

Account Number ☐☐☐☐☐☐☐☐☐☐☐☐☐☐☐☐ Expiration Date ☐☐☐☐

Authorized Signature
(Credit card orders without signature and expiration date cannot be processed.)

THANK YOU FOR YOUR ORDER

PRICES SUBJECT TO CHANGE WITHOUT PRIOR NOTICE.
PLEASE ALLOW 2 WEEKS FOR DELIVERY.
NO REFUNDS NO RETURNS.

* Data General is required by law to collect applicable sales or use tax on all purchases shipped to states where DG maintains a place of business, which covers all 50 states. Please include your local taxes when determining the total value of your order. If you are uncertain about the correct tax amount, please call 508–870–1600.

Form 702
Rev. 8/87

DATA GENERAL CORPORATION
TECHNICAL INFORMATION AND PUBLICATIONS SERVICE

TERMS AND CONDITIONS

Data General Corporation ("DGC") provides its Technical Information and Publications Service (TIPS) solely in accordance with the following terms and conditions and more specifically to the Customer signing the Educational Services TIPS Order Form. These terms and conditions apply to all orders, telephone, telex, or mail. By accepting these products the Customer accepts and agrees to be bound by these terms and conditions.

1. CUSTOMER CERTIFICATION
Customer hereby certifies that it is the owner or lessee of the DGC equipment and/or licensee/sub–licensee of the software which is the subject matter of the publication(s) ordered hereunder.

2. TAXES
Customer shall be responsible for all taxes, including taxes paid or payable by DGC for products or services supplied under this Agreement, exclusive of taxes based on DGC's net income, unless Customer provides written proof of exemption.

3. DATA AND PROPRIETARY RIGHTS
Portions of the publications and materials supplied under this Agreement are proprietary and will be so marked. Customer shall abide by such markings. DGC retains for itself exclusively all proprietary rights (including manufacturing rights) in and to all designs, engineering details and other data pertaining to the products described in such publication. Licensed software materials are provided pursuant to the terms and conditions of the Program License Agreement (PLA) between the Customer and DGC and such PLA is made a part of and incorporated into this Agreement by reference. A copyright notice on any data by itself does not constitute or evidence a publication or public disclosure.

4. LIMITED MEDIA WARRANTY
DGC warrants the CLI Macros media, provided by DGC to the Customer under this Agreement, against physical defects for a period of ninety (90) days from the date of shipment by DGC. DGC will replace defective media at no charge to you, provided it is returned postage prepaid to DGC within the ninety (90) day warranty period. This shall be your exclusive remedy and DGC's sole obligation and liability for defective media. This limited media warranty does not apply if the media has been damaged by accident, abuse or misuse.

5. DISCLAIMER OF WARRANTY
EXCEPT FOR THE LIMITED MEDIA WARRANTY NOTED ABOVE, DGC MAKES NO WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, WARRANTIES OF MERCHANTABILITY AND FITNESS FOR PARTICULAR PURPOSE ON ANY OF THE PUBLICATIONS, CLI MACROS OR MATERIALS SUPPLIED HEREUNDER.

6. LIMITATION OF LIABILITY
A. CUSTOMER AGREES THAT DGC'S LIABILITY, IF ANY, FOR DAMAGES, INCLUDING BUT NOT LIMITED TO LIABILITY ARISING OUT OF CONTRACT, NEGLIGENCE, STRICT LIABILITY IN TORT OR WARRANTY SHALL NOT EXCEED THE CHARGES PAID BY CUSTOMER FOR THE PARTICULAR PUBLICATION OR CLI MACRO INVOLVED. THIS LIMITATION OF LIABILITY SHALL NOT APPLY TO CLAIMS FOR PERSONAL INJURY CAUSED SOLELY BY DGC'S NEGLIGENCE. OTHER THAN THE CHARGES REFERENCED HEREIN, IN NO EVENT SHALL DGC BE LIABLE FOR ANY INCIDENTAL, INDIRECT, SPECIAL OR CONSEQUENTIAL DAMAGES WHATSOEVER, INCLUDING BUT NOT LIMITED TO LOST PROFITS AND DAMAGES RESULTING FROM LOSS OF USE, OR LOST DATA, OR DELIVERY DELAYS, EVEN IF DGC HAS BEEN ADVISED, KNEW OR SHOULD HAVE KNOWN OF THE POSSIBILITY THEREOF; OR FOR ANY CLAIM BY ANY THIRD PARTY.

B. ANY ACTION AGAINST DGC MUST BE COMMENCED WITHIN ONE (1) YEAR AFTER THE CAUSE OF ACTION ACCRUES.

7. GENERAL
A valid contract binding upon DGC will come into being only at the time of DGC's acceptance of the referenced Educational Services Order Form. Such contract is governed by the laws of the Commonwealth of Massachusetts, excluding its conflict of law rules. Such contract is not assignable. These terms and conditions constitute the entire agreement between the parties with respect to the subject matter hereof and supersedes all prior oral or written communications, agreements and understandings. These terms and conditions shall prevail notwithstanding any different, conflicting or additional terms and conditions which may appear on any order submitted by Customer. DGC hereby rejects all such different, conflicting, or additional terms.

8. IMPORTANT NOTICE REGARDING AOS/VS INTERNALS SERIES (ORDER #1865 & #1875)
Customer understands that information and material presented in the AOS/VS Internals Series documents may be specific to a particular revision of the product. Consequently user programs or systems based on this information and material may be revision–locked and may not function properly with prior or future revisions of the product. Therefore, Data General makes no representations as to the utility of this information and material beyond the current revision level which is the subject of the manual. Any use thereof by you or your company is at your own risk. Data General disclaims any liability arising from any such use and I and my company (Customer) hold Data General completely harmless therefrom.

Managing the
DG/UX™ System

093–701088–04

Cut here and insert in binder spine pocket