

Work Group Computing User's and Administrator's Guide

Document Number 007-3484-001

CONTRIBUTORS

Written by Pete Harbeson
Illustrated by Dany Galgani
Edited by Christina Cary
Production by Linda Rae Sande
Engineering contributions by Bob Green

© 1996, Silicon Graphics, Inc.— All Rights Reserved

The contents of this document may not be copied or duplicated in any form, in whole or in part, without the prior written permission of Silicon Graphics, Inc.

RESTRICTED RIGHTS LEGEND

Use, duplication, or disclosure of the technical data contained in this document by the Government is subject to restrictions as set forth in subdivision (c) (1) (ii) of the Rights in Technical Data and Computer Software clause at DFARS 52.227-7013 and/or in similar or successor clauses in the FAR, or in the DOD or NASA FAR Supplement. Unpublished rights reserved under the Copyright Laws of the United States. Contractor/manufacturer is Silicon Graphics, Inc., 2011 N. Shoreline Blvd., Mountain View, CA 94043-1389.

Silicon Graphics and the Silicon Graphics logo are registered trademarks and Extent File System, IRIS InSight, IRIS Networker, IRIX, and XFS are trademarks of Silicon Graphics, Inc. DynaWeb is trademark of Electronic Book Technologies. FLEXIm is a registered trademark of Globetrotter Software, Inc. Macintosh is a registered trademark of Apple Computer, Inc. NetLS is a trademark of Apollo Computer, Inc., a subsidiary of Hewlett-Packard Company. Netscape Navigator is a trademark of Netscape Communications Corp. NFS is a registered trademark of Sun Microsystems, Inc. PostScript is registered trademark of Adobe Systems, Inc. UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company, Ltd.

St Peter's Basilica image courtesy of ENEL SpA and InfoByte SpA. Disk Thrower image courtesy of Xavier Berenguer, Animatica.

Work Group Computing User's and Administrator's Guide
Document Number 007-3484-001

Contents

List of Figures ix

List of Tables xi

- 1. About Work Group Computing** 1
 - The Advantages of Work Group Computing 1
 - Added Capabilities 1
 - Sharing Information 1
 - How Work Group Computing Works 2
 - Clients and Servers 2
 - Users and Administrators 2
 - Work Group Computing Administration 3
 - Choosing an Administrator and Privileged Users 3
 - About User Privileges and the Primary User 3
 - User 3
 - Privileged User 4
 - Administrator 4
 - The Responsibilities of Privileged Users 5
 - The Network Administrator's Role 5

- Using System Administration Tools 6
 - The System Toolchest 6
 - System Manager 7
 - Disk Manager 7
 - User Manager 8
 - NFS Mount Manager 8
 - Printer Manager 8
 - Software Manager 8
 - Backup & Restore 8
 - View System Log 9
 - Run Confidence Tests 9
 - Logging In as root Through a Shell Window 9
- 2. Filesystems and Files 11**
 - Filesystem Administration 11
 - About Filesystems 11
 - Inodes 13
 - Types of Files 14
 - Hard Links and Symbolic Links 14
 - Filesystem Names 16
 - Setting Up Shared Filesystems 16
 - Network File Systems (NFS) 16
 - Maintaining Filesystems 17
 - Mounting and Unmounting Filesystems 17
 - Manually Mounting Filesystems 17
 - Mounting Filesystems Automatically With the /etc/fstab File 18
 - Mounting a Remote Filesystem Automatically 18
 - Unmounting Filesystems 18
 - Administration Considerations 19
 - Backup and Recovery 19
 - Types of Backup Media 20
 - Choosing a Backup Tool 20

Backup Strategies	21
When and What to Back Up	22
Root Filesystems	22
User Filesystems	23
Incremental Backups	23
Backing Up Files Across a Network	23
Automatic Backups	24
Storing Backups	25
How Long to Keep Backups	25
Reusing Tapes	26
System Backups	26
General Backup Procedure	27
3. Installing and Setting Up Software	29
Overview of Software Installation	29
Software Product	29
About Software Product Releases	30
Product Descriptions	30
Installation Database	30
Images	31
Patch Releases	31
How Software Installation Works	31
Making Installable Software Available to the Work Group	32
Selecting a Distribution Source	32
Setting Up an Installation Server	32
Creating a Distribution Directory	33
Configuring an Installation Account	33
Backing Up the Target Systems	34
Planning the Order of Installation	34
About Miniroot Installations	35
Installing Reference Pages	35
Installing Software Using the Software Manager	36

- System Software Maintenance 36
 - Installing a Software Update 36
 - Installing Optional Software Products 37
 - Installing Patch Releases 37
 - Installing Software for Hardware Upgrades 37
 - Installing Accompanying Product Releases 37
 - Reinstalling the Same Software 38
 - Removing All Software 38
 - Troubleshooting Software Installations 38
- 4. Software Licenses for Your Work Group 41**
 - Overview of Software Licensing and Administration 41
 - What Is a License? 42
 - How Licensing Works 42
 - License Groups for Concurrent Licensing 43
 - Licensing Software 43
 - Nodelocked License Startup 43
 - Concurrent License Startup 44
 - Planning and Setting Up Your Licensing Scheme 44
 - Choosing Whether to Implement License Groups 45
 - Setting Up License Groups 45
 - Choosing Systems to be Global Location Brokers and Network License Servers 45
 - Setting Up Global Location Brokers and Network License Servers 46
 - The Licensing Process 46
 - Maintenance 50
 - Troubleshooting Concurrent Licenses 51

5.	Identifying Your Work Group	53
	About Work Group Identification	53
	Introduction to NIS	53
	NIS Clients and Servers	54
	NIS Servers	54
	NIS Maps	54
	NIS Domains	54
	NIS and Internet Domains	55
	Setting Up NIS	56
	Setting Up the NIS Master Server	56
	Setting Up NIS Clients	57
	Administering NIS	57
	Troubleshooting	58
6.	Peripheral Devices	59
	About Peripheral Devices	59
	Setting Up Peripheral Devices	59
	Printers	60
	Configuring a Network Print Server	60
	Configuring a Print Client	61
	Printer Administration	61
	CD-ROM, Floptical, and Floppy Disk Drives	62
	CD-ROM Filesystems	62
	Floptical and Floppy Disk Filesystems	63
	Tape Drives	63
	Adding a Tape Drive	64
	Troubleshooting	64
	Troubleshooting Your Printing System	64
	Troubleshooting Tape Drives	65
	Troubleshooting Tape Read Errors	66

- 7. **Security** 67
 - Overview of System Security 67
 - Components of Security 68
 - Physical Security 68
 - Passwords 68
 - PROM Passwords 69
 - Second (Dialup) Passwords 69
 - Creating a Shadow Password File 70
 - Password Aging 70
 - Checking the Password File 70
 - File and Directory Permissions 71
 - Security Checklists 72
 - Security for Administrators 73
- 8. **Sharing Documentation** 75
 - Overview and Benefits of Sharing Documentation 75
 - IRIS InSight and DynaWeb 75
 - How DynaWeb Works 75
 - DynaWeb Directory Structure 76
 - Collections 77
 - Home Page 77
 - Publisher's Page 77
 - Access and Error Logs 78
 - CGI Scripts 78
 - Clickable Graphics 78
 - Setting Up DynaWeb for Your Work Group 78
 - DynaWeb Setup Options 78
 - DynaWeb Setup Process 79
 - Troubleshooting DynaWeb 80
 - Directories 80
 - Home Page 81
 - Client Problems 81
 - Authentication 82

List of Figures

Figure 2-1	IRIX Filesystems	12
Figure 4-1	Licensing Process for New Software (part 1)	47
Figure 4-2	Licensing Process for New Software (part 2)	48
Figure 4-3	Licensing Process for a Replacement License	49
Figure 5-1	Basic NIS Domain	55

List of Tables

Table 2-1	Filesystem File Types	14
------------------	-----------------------	----

About Work Group Computing

The Advantages of Work Group Computing

Work group computing isn't just having a workstation on every desk. It means using the capabilities of your Silicon Graphics® systems and network to make something more than a collection of individual systems: a work group system. In a very real sense, work group computing means that all of your workstations, servers, printers, drives, and other devices make up a single system. A particular workstation might lack a CD-ROM drive, or start to run out of disk space during an intensive project. But if any other workstation has a CD-ROM, or if other workstations have ample disk space available, those resources are completely available because they're part of the work group system.

Work group computing offers power, resources, communication, flexibility, and capabilities that make each workstation much more than it could be as a standalone computer.

Added Capabilities

A work group system has capabilities that no set of individual computers can offer. Examples include centralized, automatic backup of important files, access to the hard disk storage on the entire work group system, and access to peripheral devices no matter where they are physically connected. When you use a work group system, you can use the *whole* system—not just the part of it that's on your desk.

Sharing Information

In work group computing, the information on one system can be available to other systems. If one member of your work group has a report, an image, or a document that someone else needs, it's easy to gain access to it.

An important component of sharing information in a work group computing system is consistent naming. All of the computers in the work group, for example, have names.

Each computer must have a record of all of the other machine names, and of the names of the users of those machines. The Network Information System (NIS) provides an automated way of keeping records of those names, as well as a wealth of other information about all the machines and users in the work group.

How Work Group Computing Works

Work Group computing is a concept as well as a collection of hardware and software connected by a network. The concept is that the computing resources of an organization should be available at the time and place they're needed. When Jennifer needs additional disk space to store a massive 3D rendering, it doesn't make sense to purchase that storage when it already exists. It might exist somewhere else in the work group; on Ben's machine, for example. Work Group computing is primarily about designing a computing environment in which the resources of the whole system are flexible enough to meet the needs of any user, regardless of their physical location or their personal workstation's physical capabilities.

Clients and Servers

Clients and servers are integral to work group computing. Clients and servers are not physical machines; they are processes that have relationships to one another. A server is a process that provides some kind of service (for example, a database of the names of all of the machines on the network). A client is a process that makes use of a service (for example, by enabling a user to find a machine on the network). A server generally provides services to many clients. Both a client process and a server process can take place on the same machine, or on different machines. It doesn't matter which, and in fact, you often can't even tell!

Users and Administrators

Work Group computing also makes a distinction between users and administrators. This is a distinction in roles, not a distinction of people. As a member of your work group, most of the time your role is that of a user. You may be authoring content, using information, communicating, editing, or what-have-you. User roles involve activities in which the computing systems you use are tools—means to an end. Sometimes, however, you will take an administrative role. When you're backing up your files, setting up a password, or configuring a piece of software, you're an administrator, even if just for a

few minutes. Administrator roles involve activities in which the computing systems are the objects of your attention.

Everyone in a work group is an administrator some of the time, and shares responsibility for some portion of the administration of the system. With power comes responsibility; that's the responsibility that comes with the power of work group computing.

Work Group Computing Administration

Choosing an Administrator and Privileged Users

Within your own work group, it is important to decide which people will be responsible for keeping the system in good running order, and, if the system is connected to a network, who will work in conjunction with the network administrator to access network services. In some work groups, you may also need to designate a network administrator.

Administrative tasks are not, in many work group environments, allocated to just one person. Even if the work group tends to have one primary administrator, system administration is seldom that person's only job. Because administrative responsibilities are often shared in a work group, it is important that each user have some familiarity with those responsibilities.

About User Privileges and the Primary User

Because many people may use the same system, the system provides a built-in security scheme where you can grant different people different capabilities for changing the system. There are three levels of capability:

User

A user is any person who has a login account on a system. After logging in, users can change only their personal work areas. A user can run the graphical administration tools from the System toolchest, but the features of the tools that change system information are not available.

Privileged User

A person whose login account includes administrative privileges is a privileged user. When privileged users log in, they can change their personal work areas, and can use the graphical administration tools to change or customize the entire system (for example, to set up a disk or create a login account). There can be more than one privileged user on the same system.

Administrator

The administrator of a system is the person who can use the most privileged account, the *root* account. This person should have a personal login account for daily use, but, when there are serious system problems to correct, the person logs in to the root account to change system information using the graphical tools or using the IRIX™ shell.

The administrator has all the capabilities of a privileged user, as well as the ability to change information in the root account (such as the password) and to log in to an IRIX shell as root. Because there is only one root account on a system, there is only one administrator per system. The System Manager window for a particular system includes the name of the system's administrator so other users know who to contact for help.

In a work group computing environment where each person has one system, the responsibilities of users, privileged users, and administrators are often allocated in this manner:

- Each person, as the "system owner," is completely responsible for maintaining his or her own system. That person is both a privileged user and the administrator for that system. Usually the person performs administrative tasks while logged in to a personal account. But when the IRIX shell with administrative privileges is needed, the person logs out and logs back in to the root account.
- The system owner logs in as a privileged user in order to add login accounts for other people who occasionally need to use the system. If one of these users ever needs to perform administrative tasks, the system owner adds privileges to the account to make that user a privileged user.

For environments in which one person uses a particular system much more frequently than anyone else (where the person is essentially the system's owner), you can designate that person as the primary user. The primary user does not necessarily have any special access privileges, but the person's name appears along with the administrator's name in the system's System Manager window so other users know who uses the system regularly. There is only one primary user per system.

Note: The System Setup tool supports the model where one person has one system that he or she must maintain. When you create a login account for a person using this tool, it designates that person as the primary user, and also makes the person a privileged user. If necessary, you can later use the User Manager to remove administrative privileges or assign the title of primary user to someone else.

The Responsibilities of Privileged Users

A privileged user can use administration tools to perform the tasks listed below. For more information about these tasks, see the *Personal System Administration Guide*.

- Setting up the system initially as a standalone system or as a member of an existing network.
- Creating login accounts so all users of the system can access it. If the system will be connected to a network, the administrator may work in conjunction with the network administrator.
- Connecting any peripheral devices and configuring software so that the devices work properly.
- Monitoring and troubleshooting the system to keep it working efficiently and properly.
- Contacting the network administrator before connecting a system to the network. The network administrator provides information that uniquely identifies each system on the network and ensures that the regular users of a system can have accounts on other systems on the network.
- Making all, some, or none of a system's directories available to all, some, or none of the other systems and users on the network.
- Providing access to printers on other systems so the users of a system can send files to them.
- Providing access to disk space that's available on other systems on the network.

The Network Administrator's Role

The responsibilities of a network administrator vary greatly from site to site. In a work group environment, the network administrator is often responsible for these tasks:

- Setting up and maintaining the network so connections are reliable and data is transferred as quickly as possible.
- Creating, maintaining, and periodically distributing a list of all systems and users so that each has a unique identity on the network.
- Setting up and maintaining network services such as electronic mail and the Network Information Services (NIS).

Using System Administration Tools

As the administrator or privileged user you can use two different types of tools:

- The System toolchest provides a collection of graphical system administration tools.
- The IRIX shell accepts IRIX commands that you use for more advanced administrative tasks.

This online information describes how to use the graphical tools in the System toolchest to perform as many administration tasks as possible; in cases where no graphical tools support a task, you must use IRIX commands or edit system files. If you prefer to perform all administrative tasks without using the System toolchest, see the *IRIX Admin* manual set (choose “Online Books” from the System toolchest, and look in the SGI_Admin bookshelf). Regardless of whether you edit system files manually or let the graphical tools do it for you, you are changing the same system files.

The System Toolchest

The administrative tools that you’ll use most frequently are in the System toolchest menu; you run a tool by choosing it from the menu. The first tool in the menu is the *System Manager*. Tools that you use less frequently appear when you choose “System Admin Tools” from the Tools menu in the System Manager window.

The System toolchest includes these tools:

- System Manager
- Disk Manager
- User Manager
- NFS Mount Manager

- Printer Manager
- Software Manager
- Backup & Restore
- View System Log
- Run Confidence Tests

Each tool is described in a section below.

System Manager

Using the System Manager, a user can perform these tasks::

- Determine which of the system's resources are available for use by other systems on the network.
- Check the hardware and software inventory of the system.
- Get business card information about the system's primary user and administrator.

In addition, a privileged user can perform these tasks:

- Run all the system administration tools
- Designate the primary user.

Disk Manager

The Disk Manager shows how much disk space you are using and how much is still available. Privileged users can specify when the system should warn that the disk is nearly full. When you install a new disk, a privileged user also uses this tool to specify a directory from which to access the disk (the directory from which you access a disk is called the *mount point* for that disk).

User Manager

The User Manager enables users to perform these tasks:

- View information about their own login account.
- View information about other login accounts.
- Change business card information about their own account.

In addition, privileged users can use the User Manager to create, change, and delete user login accounts.

NFS Mount Manager

The NFS[®] Mount Manager is available only on systems that have the optional NFS software installed. The NFS Mount Manager lets privileged users access (mount) directories on other systems so that users can access the directories as if they resided on the local system.

Printer Manager

With the Printer Manager, users can view the list of available printers and drag printer icons onto their desktops for use. Additionally, privileged users can use the Printer Manager to set up the software for local or remote printers so the system can access the printers.

Software Manager

Users and privileged users can view a list of installed software by using the Software Manager. The Software Manager also enables the administrator to install and remove software.

Backup & Restore

Backup & Restore enables users to to back up their own files and directories. The administrator can use Backup & Restore to create full system backups from which the entire system can be restored in the event of a serious system failure.

View System Log

View System Log displays a log of all the system messages, and lets a privileged user customize when and how users should be notified of system problems.

Run Confidence Tests

Run Confidence Tests displays a window in which users can test various peripherals and system components, such as the mouse, keyboard, and monitor.

Logging In as root Through a Shell Window

Only the administrator can perform administrative tasks that are not supported by the graphical tools since these tasks require the use of the root account in a shell window. The home directory for the root account is the root (/) directory of the filesystem. The user logged in to the root account can move, change, and delete every file and directory on the system, regardless of who owns them and what type of permissions they have set. Be sure to create a password for this account that only the administrator knows.

Note: Some UNIX® and IRIX documents refer to the user of the root account as the *superuser* rather than the administrator.

When you're already logged in as a regular user, you can start a shell window and log in as root by following these steps:

1. Choose "Unix Shell" from the Desktop toolchest.
2. Position your cursor within the new window and enter:

```
% login root
```

If a prompt for a password appears, type the password, then press Enter. If a prompt appears but the root account has no password, just press Enter. (See "Customizing System Account Information" to create, change, or remove a password.)

You are now logged in to the root account and are located in the root (/) directory. When you are logged in as root, the IRIX prompt is a pound sign (#) rather than a percent sign (%).

To log out of the root account, enter:

```
# logout
```

The shell window disappears.

Filesystems and Files

Filesystem Administration

Setting up and maintaining a filesystem is a system administrative task. If you are primarily a user in your work group, read the first part of this chapter to get a general understanding of filesystems, and read the section “Backup and Recovery” so that you understand what you need to do to protect the information on your system. If you are responsible for system administration in your work group, read all of this chapter to fully understand filesystems, including setup, maintenance, backup, and recovery.

About Filesystems

If you store paper documents in a filing cabinet, you can choose to organize them any number of ways (for example, alphabetically, by number, or by date). The method you choose is your filing system. In much the same way, when you store information on a disk connected to your workstation, the manner in which it is stored is determined by the installed *filesystem*. Your Silicon Graphics system can use any of several different filesystems.

Technically speaking, a filesystem is a data structure that organizes files and directories on a disk partition so that they can be easily retrieved. Only one filesystem can reside on a disk *partition* (a disk always has at least one partition, and may have several).

A file is a one-dimensional array of bytes that can be in any possible order and can represent literally any information, from a memo to a three-dimensional model to an application program. Information about each file is stored in structures called *inodes* (inodes are described in the next section, “Inodes”). A file exists within a single filesystem.

A directory is a container that stores files and other directories. It is merely another type of file that the user is permitted to use, but not allowed to write; the operating system itself retains the responsibility for writing directories. Directories cannot span filesystems. The combination of directories and files make up a filesystem.

The starting point of any filesystem is an unnamed directory that serves as the root for that particular filesystem. In the IRIX operating system there is always one filesystem that is itself referred to by that name, the *root* filesystem. Traditionally, the root directory of the root filesystem is represented by a single slash (/). Filesystems are attached to the directory hierarchy by the *mount* command. The following illustration shows root filesystem with a mount point of a single slash (/), and two additional filesystems with mount points of /d2 and /d3, respectively

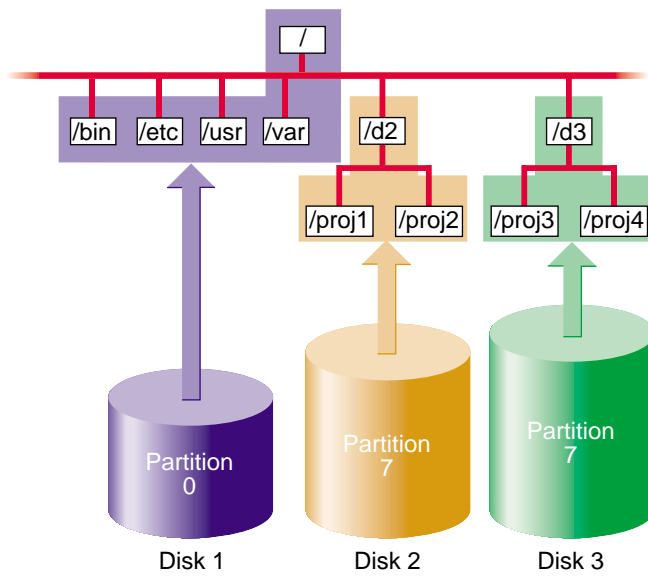


Figure 2-1 IRIX Filesystems

You can join two or more disk partitions to create a *logical volume*. The logical volume can be treated as if it were a single disk partition. A filesystem can reside on a logical volume, thus making it possible for a single filesystem to exist across more than one disk. You might create a logical volume with a single filesystem if you need to deal with very large files, or if you need to make it convenient to access a very large amount of storage space.

The following subsections describe key components of filesystems in more detail. For even greater detail, see *IRIX Admin: Disks and Filesystems*.

Inodes

Information about each file is stored in a structure called an inode. The word inode is an abbreviation of the term *index node*. An inode is a data structure that stores all information about a file except its name, which is stored in the directory. Each inode has an identifying inode number, which is unique across the filesystem that includes the file.

An inode contains this information:

- the type of the file (see the next section, “Types of Files,” for more information)
- the access mode of the file; the mode defines the access permissions read, write, and execute and may also contain security labels and access control lists
- the number of hard links to the file
- who owns the file (the owner’s user-ID number) and the group to which the file belongs (the group-ID number)
- the size of the file in bytes
- the date and time the file was last accessed and last modified
- information for finding the file’s data within the disk partition or logical volume
- the pathname of symbolic links (when they fit and on XFS filesystems only)

You can use the `ls` command with various options to display the information stored in inodes. For example, the command `ls -l` displays all but the last two items in the list above in the order listed (the date shown is the last modified time).

Inodes do not contain the name of the file or its directory.

Types of Files

Filesystems can contain the file types listed in Table 2-1. You can see the type of a file in a shell window when you use the `ls -l` command—the first character in each line of output indicates the file’s type.

Table 2-1 Filesystem File Types

File Type	Indicator	Comments
Regular files	-	One-dimensional arrays of bytes.
Directories	d	Containers for files and other directories.
Symbolic links	l	Files that contain the name of another file or a directory.
Character	c	Enable communication devices between hardware and IRIX; data is accessed on a character-by-character basis.
Block Device	b	Enable communication between hardware and IRIX; data is accessed in blocks from a system buffer cache.
Named pipe (also known as FIFOs)		Allow communication between two unrelated processes on the same host. They are created with the <code>mknod</code> command.
UNIX domain sockets	s	Connections between processes that allow them to communicate, possibly over a network.

Hard Links and Symbolic Links

As discussed in the section “Inodes” in this chapter, information about each file, except for the name and directory of the file, is stored in an inode for the file. The name of the file is stored in the file’s directory and a link to the file is created by associating the filename with an inode number. This type of link is called a *hard link*. Although every file is a hard link, the term is usually used only when two or more filenames are associated with the same inode number. Because inode numbers are unique only within a filesystem, hard links cannot be created across filesystem boundaries.

The second and later hard links to a file are created with the `ln` command, without the `-s` option. For example, say the current directory contains a file called *origfile*. To create a hard link called *linkfile* to the file *origfile*, enter this command:

```
% ln origfile linkfile
```

The output of `ls -l` for *origfile* and *linkfile* shows identical sizes and last modification times:

```
% ls -l origfile linkfile
-rw-rw-r--  2 joyce  user      4 Apr  5 11:15 origfile
-rw-rw-r--  2 joyce  user      4 Apr  5 11:15 linkfile
```

Because *origfile* and *linkfile* are simply two names for the same file, changes in the contents of the file are visible when using either filename. Removing one of the links has no effect on the other. The file is not removed until there are no links to it (the number of links to the file, the link count, is stored in the file's inode).

Another type of link is the *symbolic link*. This type of link is actually a file. The file contains a text string, which is the pathname of another file or directory. Because a symbolic link is a file, it has its own owners and permissions. The file or directory it points to can be in another filesystem. If the file or directory that a symbolic link points to is removed, it is no longer available and the symbolic link becomes useless until the target is recreated (it is called a *dangling symbolic link*).

Symbolic links are created with the `ln` command with the `-s` option. For example, to create a symbolic link called *linkdir* to the directory *origdir*, enter this command:

```
% ln -s origdir linkdir
```

The output of `ls -ld` for the symbolic link is shown below. Notice that the permissions and other information don't match. The listing for *linkdir* shows that it is a symbolic link to *origdir*.

```
% ls -ld linkdir origdir
drwxrwxrwt 13 sys  sys 2048 Apr  5 11:37 origdir
lrwxrwxr-x  1 joyce  user   8 Apr  5 11:52 linkdir -> origdir
```

When you use “`..`” in pathnames that involve symbolic links, be aware that “`..`” refers to the parent directory of the true file or directory, not the parent of the directory that contains the symbolic link.

Filesystem Names

Filesystems don't have names per se; they are identified by their location on a disk or their position in the directory structure in these ways:

- by the block and character device filenames of the disk partition or logical volume that contains the filesystem
- by a mnemonic name for the disk partition or logical volume that contains the filesystem
- by the mount point for the filesystem

The filesystem identifier from the list above that you use with commands that administer filesystems (such as *mkfs*, *mount*, *umount*, and *fsck*) depends upon the command. See the reference page for the command you want to use or examples in this guide to determine which filesystem name to use.

Setting Up Shared Filesystems

Network File Systems (NFS)

NFS filesystems are available if you are using the optional NFS software. NFS filesystems are exported from one host and mounted on other hosts across a network. The capabilities of NFS filesystems meet the needs of many work groups, and are often used in work group computing environments.

On the hosts where the filesystems reside, they are treated just like any other filesystem (for example the Extent File System™—EFS—or XFS). The only special feature of these filesystems is that they can be *exported* and mounted on other workstations. Exporting NFS filesystems is done with the *exportfs* command. On other hosts, these filesystems are mounted with the *mount* command or by using the automount facility of NFS.

NFS filesystems are described in detail in the *ONC3/NFS Administrator's Guide*, which is included with the NFS software option.

Maintaining Filesystems

Mounting and Unmounting Filesystems

Filesystems must be mounted in order to be accessed by IRIX. The subsections below explain the use of the *mount* and *umount* commands and the file */etc/fstab* to mount and unmount filesystems.

Tip: You can mount and unmount XFS filesystems using the graphical user interface of the *xfsm* command. For information, see its online help.

Manually Mounting Filesystems

The *mount* command is used to manually mount filesystems. The basic forms of the mount command are as shown below:

```
% mount device_file mount_point_directory
% mount host:directory mount_point_directory
```

device_file is a block device file. *host:directory* is the hostname and pathname of a remote directory that has been exported on the remote host by using the *exportfs* command on the remote host (it requires NFS). *mount_point_directory* is the mount point directory. The mount point must already exist (you can create it with the *mkdir* command).

If you omit either the *device_file* or the *mount_point_directory* from the *mount* command line, *mount* checks the file */etc/fstab* to find the missing argument.

For example, to mount a filesystem manually, use this command:

```
% mount /dev/dsk/dks0d1s6 /usr
```

Another example, which uses a mnemonic device file name, is this:

```
% mount /dev/usr /usr
```

See the *mount* reference page for more information about the *mount* command.

Mounting Filesystems Automatically With the */etc/fstab* File

The */etc/fstab* file contains information about every filesystem and swap partition that is to be mounted automatically when the system starts up. In addition, the */etc/fstab* file is used by the *mount* command when only the device block file or the mount point is given to the *mount* command.

For details about adding entries to */etc/fstab*, and for information about the format of the */etc/fstab* file, see the manual *IRIX Admin: Disks and Filesystems*.

Mounting a Remote Filesystem Automatically

If you have the optional NFS software, you can automatically mount any remote filesystem whenever it is accessed (for example, by changing directories to the filesystem with *cd*). The remote filesystem must first be exported with the *exportfs* command.

For complete information about setting up automounting, including all the available options, see the automount and exportfs reference pages. These commands are discussed more completely in the *ONC3/NFS Administrator's Guide* .

Unmounting Filesystems

Filesystems are automatically unmounted when the system is shut down. To manually unmount filesystems, use the *umount* command. There are three basic forms of the command, as shown below. You can use the first two forms to unmount local filesystems, and the first and third forms to unmount remote filesystems.

```
% umount mount_point_directory
% umount device_file
% umount host:directory
```

For example, to unmount a local or remote filesystem mounted at */d2*, enter this command:

```
% umount /d2
```

To unmount the filesystem on the partition */dev/dsk/dks0d1s7*, enter this command:

```
% umount /dev/dsk/dks0d1s7
```

To unmount the remote-mounted (NFS) filesystem *depot:/usr/spool/news*, enter this command:

```
% umount depot:/usr/spool/news
```

To be unmounted, a filesystem must not be in use. If it is in use and you try to unmount it, you get a “Resource busy” message. These messages and their solutions are explained in the *umount* reference page.

Administration Considerations

To administer filesystems, you need to do the following:

- Monitor the amount of free space and free inodes available.
- If a filesystem is chronically short of free space, take steps to alleviate the problem, such as removing old files and imposing disk usage quotas.
- If you use EFS filesystems, use *fsck* periodically to check the filesystems for data integrity. (XFS filesystems are checked with *xfs_check* only when a problem is suspected.)
- Back up filesystems.

You can perform many of these functions automatically by using *shell scripts*. Many shell scripts are installed on your system by default. For more information, examine the scripts in */usr/lib/acct*, such as *ckpacct* and *remove*, for ideas about how to build your own administration scripts.

Backup and Recovery

No matter what kind of computing systems and environment you use, it is essential to regularly back up files. Backup copies are often the only way to recover files that have been accidentally erased, or lost due to hardware problems. It is the responsibility of the work group administrator to make sure there are backups of the files used in the work group. Depending on the work group’s backup plan, it may be up to users or to the administrator to actually perform the backups, as well as any recovery that might be needed.

Types of Backup Media

Some of the common types of backup media supported on Silicon Graphics systems include the following:

- 1/4" cartridge tape, 4-track
- 8 mm cartridge
- DAT
- DLT

In addition to backup devices attached to any particular system, backup devices of various types and capacities may be accessible through network connections. Refer to your owner's guide for information on locally accessible devices, and the appropriate vendor documentation for network-accessible device information.

Choosing a Backup Tool

The IRIX system provides a variety of backup tools, and you should use whichever ones work best for you. If many users at your site are already familiar with one backup program, you may wish to use that program consistently. If there are workstations at your site from other manufacturers, you may wish to use a backup utility that is common to all the workstations.

IRIX provides the following utilities for backing up your data:

- System Manager, providing backup functions (this is the recommended tool if you are backing up your own filesystem).
- *bru*, a backup tool featuring automatic file compression, space estimates, and integrity checking.
- *Backup* and *Restore*, command-line front ends to the *bru* utility.
- *dump* and *restore*, standard UNIX utilities (cannot back up XFS filesystems).
- *xfsdump* and *xfrestore* for XFS filesystems, XFS compatible versions of *dump* and *restore*.
- *tar*, the most common UNIX backup utility (almost certainly available on workstations from other manufacturers).

- *cpio*, a standard UNIX command (often combined in command line pipes with other commands).
- *dd*, a standard UNIX command to read input and write output.

Optional products for Silicon Graphics systems are also available. IRIS NetWorker™ is a scalable, full-featured data management tool for data backup and recovery. You can use IRIS NetWorker to back up data on high-end servers, or centrally manage backups for all your network workstations and file servers.

For more information about backup tools, see the *Personal System Administration Guide* and *IRIX Admin: Backup, Security, and Accounting*.

Backup Strategies

As the work group administrator, you should develop a regimen for backing up the system or systems at your site and follow it closely. That way, you can accurately assess which data you can and cannot recover in the event of a mishap.

Exactly how you perform backups depends upon your workstation configuration and other factors. Regardless of the strategy you choose, though, you should always keep at least two full sets of reasonably current backups. You should also encourage users to make their own backups, particularly of critical, rapidly changing files. Users' needs can change overnight, and they know best the value of their data.

Workstation users can back up important files using the System Manager, found in the System toolchest on your screen. The System Manager is described in detail in the *Personal System Administration Guide*. Make sure users have access to an adequate supply of media (for example, cartridge tapes), whether new or used.

If your media can handle your largest filesystem with a single volume, you don't have to use an incremental backup scheme, though such a system reduces the amount of time you spend making backups. However, if you must regularly use multiple volumes to back up your filesystems, then an incremental backup system reduces the number of tapes you use.

The following sections discuss the different aspects of backing up data.

When and What to Back Up

How often you back up your data depends upon how busy a system is and how critical the data is. A simple rule of thumb is to back up any data on the system that is irreplaceable or that someone does not want to reenter.

Root Filesystems

On most systems, the *root* filesystem is fairly static. You do not need to back it up as frequently as the */usr* filesystem.

Changes may occur when you add software, reconfigure hardware, change the site-networking (and the system is a server or network information service [NIS] master workstation), or change some aspect of the workstation configuration. In some cases, you can maintain backups only of the individual files that change, for example, */unix*, */etc/passwd*, and so forth.

This process of backing up single files is not always simple. Even a minor system change such as adding a user affects files all over the system, and if you use the graphical System Manager, you may tend to forget all the files that may have changed. Also, if you are not the only administrator at the site, you may not be aware of changes made by your coworkers. Using complete filesystem backup utilities, such as the System Manager or *bru*, on a regular schedule avoids these problems.

A reasonable approach is to back up the root partition once a month. In addition to regular backups, here are some specific times to back up a root filesystem:

- whenever you add users to the system, especially if the workstation is an NIS master workstation
- just before installing new software
- after installing new software and when you are certain the software is working properly

If your system is very active, or if you are not the only administrator, back up the root filesystem regularly.

User Filesystems

The */usr* filesystem, which often contains both system programs (such as in */usr/bin*) and user accounts, is usually more active than a root filesystem. Therefore, you should back it up more frequently.

At a typical multiuser installation, backing up once per day, using an incremental scheme, should be sufficient.

Treat the */var* filesystem similarly—it contains data such as the contents of users' mailboxes.

Incremental Backups

Incremental backups can use fewer tapes to provide the same level of protection as repeatedly backing up the entire filesystem. They are also faster than backing up every file on the system.

An incremental scheme for a particular filesystem looks something like this:

1. On the first day, back up the entire filesystem. This is a monthly backup.
2. On the second through seventh days, back up only the files that changed from the previous day. These are daily backups.
3. On the eighth day, back up all the files that changed the previous week. This is a weekly backup.
4. Repeat steps 2 and 3 for four weeks (about one month).
5. After four weeks (about a month), start over, repeating steps 1 through 4.

You can recycle daily tapes every month, or whenever you feel safe about doing so. You can keep the weekly tapes for a few months. You should keep the monthly tapes for about one year before recycling them.

Backing Up Files Across a Network

If you are managing a site with many networked workstations, you may wish to save backups on a device located on a central workstation.

To back up across a network, use the same basic backup commands, but with a slight change. Enter:

```
% system_name:/dev/tape
```

If required, specify an account on the remote device:

```
% user@system_name:/dev/tape
```

Users can use a central tape drive from their workstations with this method. Note that if you are backing up to a remote tape drive on a workstation that is not made by Silicon Graphics, the device name */dev/tape* may not be the correct name for the tape drive. Always learn the pathname of the tape device before executing the backup commands.

For example:

```
% tar cvf guest@alice:/dev/tape ./bus.schedule
```

or

```
% echo "./bus.schedule" | cpio -ovc0 guest@alice:/dev/tape
```

Automatic Backups

You can use the *cron* utility to automatically back up filesystems at predetermined times. The backup media must be already mounted in the drive, and, if you want this to be truly automatic, it should have enough capacity to store all the data being backed up on a single piece of media. If all the data doesn't fit, then someone must manually change backup media.

Here is an example *cron* command to back up the */usr/src* hierarchy to */dev/tape* (tape drive) every morning at 03:00 using *bru*:

```
0 3 * * * /usr/sbin/bru -c -f /dev/tape /usr/src
```

Place this line in a crontabs file, such as */var/spool/cron/crontabs/root*.

This sort of command is useful as a safety net, but you should not rely solely on automatic backups. There is no substitute for having a person monitor the backup process from start to finish and properly archive and label the media when the backup is finished. For more information on using *cron* to perform jobs automatically, see *IRIX Admin: System Configuration and Operation*.

Storing Backups

Store your backup tapes carefully. Even if you create backups on more durable media, such as optical disks, take care not to abuse them. Set the write protect switch on tapes you plan to store as soon as a tape is written, but remember to unset it when you are ready to overwrite a previously-used tape.

Do not subject backups to extremes of temperature and humidity, and keep tapes away from strong electromagnetic fields. If there are a large number of workstations at your site, you may wish to devote a special room to storing backups.

Store magnetic tapes, including 1/4" and 8 mm cartridges, upright. Do not store tapes on their sides, as this can deform the tape material and cause the tapes to read incorrectly.

Make sure the media is clearly labeled and, if applicable, write-protected. Choose a label-color scheme to identify such aspects of the backup as what system it is from, what level of backup (complete versus partial), what filesystem, and so forth.

To minimize the impact of a disaster at your site, such as a fire, you may want to store main copies of backups in a different building from the actual workstations. You have to balance this practice, though, with the need to have backups handy for recovering files.

If backups contain sensitive data, take the appropriate security precautions, such as placing them in a locked, secure room. Anyone can read a backup tape on a system that has the appropriate utilities.

How Long to Keep Backups

You can keep backups as long as you think you need to. In practice, few sites keep system backup tapes longer than about a year before recycling the tape for new backups. Usually, data for specific purposes and projects is backed up at specific project milestones (for example, when a project is started or finished). As site administrator, you should consult with your users to determine how long to keep filesystem backups.

With magnetic tapes, however, there are certain physical limitations. Tape gradually loses its flux (magnetism) over time. After about two years, tape can start to lose data.

For long-term storage, re-copy magnetic tapes every year to year-and-a-half to prevent data loss through deterioration. When possible, use checksum programs, such as the *sum* utility, to make sure data hasn't deteriorated or altered in the copying process. If you want to store data reliably for several years, consider using optical disk.

Reusing Tapes

You can reuse tapes, but with wear, the quality of a tape degrades. The more important the data, the more precautions you should take, including using new tapes.

If a tape goes bad, mark it as “bad” and discard it. Write “bad” on the tape case before you throw it out so that someone doesn’t accidentally try to use it. Never try to reuse an obviously bad tape. The cost of a new tape is minimal compared to the value of the data you are storing on it.

System Backups

To make a backup of your system on any system with a graphical user interface, bring up the System menu on the System Toolchest and select the Backup & Restore menu item. From the Backup & Restore window, follow the prompts to perform your backup. A complete set of instructions for this procedure is available in the *Personal System Administration Guide*.

Backups made with the Backup & Restore window are the easiest to make and use, and are accessible from the Recover System option on the System Maintenance Menu if they are full system backups. When you make a full system backup, the command also makes a backup of the files in the disk volume header and saves the information in a file that is stored on tape. This file is used during system recovery to restore a damaged volume header.

To make a backup of your system using an IRIX command, use the *Backup* command. Although the *Backup* command is a front-end interface to the *bru* command, *Backup* also writes the disk volume header on the tape so that the Recover System option can reconstruct the boot blocks, which are not written to the tape using other backup commands.

General Backup Procedure

Follow these steps when making a backup, no matter which backup utility you use:

1. Make sure the tape drive is clean. The hardware manual that came with your drive should state how, and how often, to clean the drive.

Dirty tape heads can cause read and write errors. New tapes shed more oxide than older tapes, so you should clean your drive more frequently if you use a lot of new tapes.

2. Make sure you have enough backup media on hand. The *bru* utility has an option to check the size of a backup, so you can determine if you have enough media. You can also use such utilities as *du* and *df* to determine the size of directories and filesystems, respectively.

Also, use good-quality media. Considering the value of your data, use the best quality media you can afford.

3. Run *fsck* first on EFS filesystems (if you are backing up an entire filesystem) to make sure you do not create a tape of a damaged filesystem. You must unmount a filesystem before checking it with *fsck*, so plan your backup schedule accordingly.

This step is not necessary if you are backing up only a few files (for example, with *tar*).

4. The default tape device for any drives you may have is */dev/tape*. If you do not use the default device, you must specify a device in your backup command line.
5. Label your backups. If you plan to reuse the media, use pencil. Include the date, time, name of the system, the name of the utility, the exact command line used to make the backup (so you'll remember how to extract the files later), and a general indication of the contents. If more than one administrator performs backups at your site, include your name.
6. Verify the backup when you are finished. Some utilities (such as *dump* and *bru*) provide explicit options to verify a backup. With other programs, you can simply list the contents of the archive—this is usually sufficient to catch errors in the backup.
7. Write-protect your media after you make the backup.
8. Note the number of times you use each tape. Keep a running tally on the tape label.

Installing and Setting Up Software

Overview of Software Installation

In a work group environment, users typically install the software they need on their own workstations. They obtain software from a central distribution source that is typically set up and maintained by the administrator. The administrator makes software available for the workgroup, and notifies the work group when additional software, updates, and patch releases become available.

Silicon Graphics software is distributed on compact discs (CDs). A CD contains one or more software products and any special tools that the products require for installation. The purpose of the installation utility, *inst*, is to transfer distribution software, which has an encoded format, to a hard disk in a format that is usable. The installation utility offers two user interfaces: a graphical interface, called Software Manager; and a command-line interface, called Inst.

The media containing installable software that you purchase from Silicon Graphics is referred to as a distribution. Most distributions are not specific to a particular model of Silicon Graphics computer; distributions ordinarily contain all versions of any hardware-specific files that might be required. Sometimes, however, a new server or personal workstation model is introduced, accompanied by a special software distribution. When this occurs, subsequent distributions are fully compatible with the new model.

Software Product

A software product is a collection of files that support either a system function or a specific application. Some products support critical functions and must be installed if the system is to operate; other products are not critical but optimize system operation and are recommended for installation by the manufacturer. A subset of required and recommended products is installed in servers and workstations before they leave the factory.

About Software Product Releases

A software product release contains all software required to support a given version of a product and the tools that are needed to install the release. One or more software product releases are contained in a distribution.

When you install a software product release, files from previous versions of the release are automatically removed before the new files are installed (the exception is configuration files, which are saved if they contain local modifications). If a product release includes hardware-specific files, the installation utility automatically determines the file version that is needed on a particular model and installs that version.

Product releases may have prerequisites that require a particular installation order. They might also have compatibility requirements with other installed products. Inst protects users against potential problems by preventing installation if there are any unresolved incompatibilities or prerequisites (unless you override this safeguard).

Installation instructions, including prerequisites and incompatibility information, are provided in the product release notes, which are always included as an installable subsystem of the product. You can read product release notes from the distribution CD using the *CDgrelnotes* or *CDrelnotes* command. Instructions for reading release notes are included in the CD “jewel case” insert.

Product Descriptions

The product description is a file that contains information about product requirements and the installation environment for the product. The name of the product description file is the short name for the product. For example, the name of the product description file for the Fortran 77 Compiler product is *ftn_dev*.

Installation Database

The installation database is a file that contains installation information for every file in the product. The name of the installation database is the shortname with an *.idb* extension. For example, the name of the installation database for the Fortran 77 Compiler product is *ftn_dev.idb*.

Images

An image is a collection of installable files that perform a similar or complementary function. Software products usually contain at least two images. Providing more than one image makes it possible to install only the software you need. For example, it is possible to install the images that contain the executable programs of a product without installing the release notes image.

Patch Releases

A patch is a collection of one or more files that correct flaws in the performance, reliability, or security of a specific software product; a patch release is a distribution containing one or more patches. Each patch remedies a specific set of reported errors in the version of the product to which the patch applies.

Every patch is identified by a unique number, such as *patchSG1234567*, for example. The number is used in manufacturing to identify the collection of errors that the patch corrects. Typically, the reported errors that result in a patch release occur only under certain operating conditions. For this reason, installing a patch is necessary only if a system or site is experiencing a problem that the patch addresses.

How Software Installation Works

Most software installations can be performed without interrupting system operations. Installations that are performed without shutting down the system are referred to as live installations. Live installations are preferred because they are usually less time-consuming and because other system operations can be sustained during the installation session.

Whenever software installation affects fundamental IRIX functions (such as device management), software must be installed by a miniroot installation, which requires a system shutdown.

Making Installable Software Available to the Work Group

The location of a software distribution is known as the distribution source; the system receiving software during an installation is known as the target. A distribution source may be a CD that is mounted on the target, a CD that is mounted on a remote system, or a centralized directory on the network to which the distribution has been copied. The directory on a distribution CD that contains the software is always called */dist*.

You will often want to set up an installation server for your work group. An installation server is a server or workstation that supplies a distribution source to other systems over the network. On your installation server, installable software is generally kept in a distribution directory, which may contain software from several distributions.

It is important to provide your work group with a consistent location for installable software. If you provide software installation services to work group members, having a consistent location makes your job easier. If work group members perform their own software installations, having a consistent location means they can be more self-sufficient.

Selecting a Distribution Source

When selecting a distribution source, consider the speed and reliability of your network, the frequency with which installations are performed, and the amount of software that will be installed. If disk space is available and the network is fast and reliable, consider creating a centralized distribution directory on an installation server. A centralized directory is particularly useful if users perform their own installations, or if the availability of the server's CD-ROM drive is subject to interruption.

When you are installing software on one or two target systems and the targets contain local CD-ROM drives, using a locally mounted distribution CD is often the most efficient distribution source, particularly if your network is slow and you plan to install a considerable amount of software. For any target that is not equipped with a CD-ROM drive, the distribution source must be a remote CD-ROM drive or distribution directory.

Setting Up an Installation Server

You can create an installation server on almost any system in your network that is reliable and has adequate disk resources. The distribution source may be a local CD-ROM drive or a distribution directory.

Any system that you plan to use as an installation server must be accessible from the remote targets. This means that communications between the server and targets must support forwarding of boot files and Trivial File Transfer Protocol (TFTP) file transfers. In addition, the installation server must contain a user account that is available to target systems.

Creating a Distribution Directory

After you have chosen an installation server, the next step is to create a distribution directory on it. The CD-ROM drive from which you copy the distribution software may be either a local or remote drive. You can copy several CDs to the distribution directory if you wish; however, do not include more than one release of a given product in the directory—different distribution directories must be created for different releases of a product.

You can create a software distribution directory that contains fewer products than are in the CD-ROM distribution by copying the files for just the products that you want. Remember that distribution directories and CD-ROM distributions have an identical structure.

For more information about creating a distribution directory, see *IRIX Admin: Software Installation and Licensing*.

Configuring an Installation Account

Installations generally use the guest account on the server. This means that the installation server's guest account must not be password protected. If the guest account on the server is either unavailable or password protected, you must provide an alternate means for accessing the server. You can allow access to the server in any of these ways:

- Remove the password from guest while installations are taking place.
- Use an account other than guest on the server (the alternate account must not be password protected) and specify the alternate account when starting the installation.
- Use a password-protected account on the server for installations and create a *.rhosts* file for the installation account. The installation account must have read permissions on the distribution source.

Backing Up the Target Systems

Although backing up the target is not a requirement for installation, it is strongly recommended. The *Personal System Administration Guide* describes general backup procedures and using System Manager to perform backups.

Consider backing up these files:

- User files that are created or copied to the target. Any file on the target that was not put there during the software installation process is considered a user file.
- Configuration files that contain information unique to the target system or the site. These files are created during installation but are likely to be modified after they are installed. The unique information in these files is not destroyed during an installation. However, the pre-installation copy of these configuration files is helpful if you decide to go back to the earlier software release after installation.

Planning the Order of Installation

If you expect to install software from more than one CD or distribution directory, you must plan the installation order, since some products require that other products be installed first.

If you are installing from multiple CDs, use the sequence numbers on their labels to put the CDs in order. Install the CDs starting with the lowest sequence number first. Use these guidelines to plan the order of your installation:

- If any CD or distribution directory that you install contains installation tools, be sure to install that CD first. CDs containing installation tools are clearly marked.
- If any CD has two sequence numbers, that CD is used twice during the installation. If you find no intervening sequence numbers, you need to insert the CD only once during the installation.
- If you have two or more CDs with the same sequence number, the order of those CDs relative to each other does not matter. For example, assume that you have four CDs with sequence numbers 400, 500, 600, and 600. The CD labeled 400 is installed first, followed by the CD labeled 500. The order of the two CDs labeled 600 is irrelevant, as long as they are installed last.

If you are installing from several distribution directories, check to see what products are in each directory. Plan to install the products in this order:

- installation tools
- operating system software
- communications software
- compilers
- optional software

About Miniroot Installations

For installations where fundamental IRIX services, such as filesystem management, are either unavailable or unreliable, a special installation method is required. This method, known as a *miniroot* installation, relies on services in the target's programmable read-only memory (PROM) to transfer special installation tools, including *Inst*, from the distribution source. This transfer is referred to as *loading the miniroot*.

For more information about miniroot installations, see *IRIX Admin: Software Installation and Licensing*.

Installing Reference Pages

The reference pages (manual pages) that apply to a product are shipped as a software component of the product.

The software subsystems in a product usually have a corresponding reference page subsystem. The names of the software and reference page subsystems differ only in the image segment of the name. The name for a reference page subsystem always contains the letters *.man* in the image segment. For example, *dmedia_dev* contains the software subsystem *dmedia_dev.sw.movie* and a reference page subsystem called *dmedia_dev.man.movie*.

Installing Software Using the Software Manager

The Software Manager lets you install and remove both operating system software and optional product software. Any user or privileged user can use the Software Manager to view a list of installed software, but only the administrator (or anyone who knows the password for the root account) can use it to install or remove software.

To install software, you specify where the new software is located. Then you can install either a default set of software by clicking *Install Automatically*, or a customized set of software by clicking *Customize Installation*.

When you click the *Customize Installation* button, you're requesting a more informative view of the software that's available. If a checkmark appeared in the "Upgrade Products" and/or "New Products" check boxes in the *Install Automatically* view, these products and their default subsystems remain selected for installation.

Note: If the root account on a system has a password, that password must be entered in order to install software.

System Software Maintenance

This section contains procedures that may be necessary to keep the software installation at your site up to date.

Installing a Software Update

When you receive a software update, it might be delivered on multiple CDs. Use the sequence numbers on the CD labels or the directions in "Planning the Order of Installation" on page 34 to determine the order in which you should install the updates. Updates are not necessarily included for all products, since not all products are updated at the same time.

When you install a software update, replace or remove all of the older subsystems in each product that you install. Do not remove just some of the older subsystems. When a server or workstation contains software subsystems from different releases of a product, compatibility problems result that can be very difficult to diagnose.

Installing Optional Software Products

Software options are products that you may or may not choose to run on a system. They are usually purchased separately from a computer system, but may also be shipped as a complimentary offering with a new system or when a new version of a purchased option is released. When you install a software option, keep these points in mind:

- The release notes for a software option explain whether a miniroot installation is required.
- If the software option requires miniroot installation, use the installation tools that accompanied the version of *oe* that is already installed on your system.
- If you have several software options to install and they are on more than one CD or distribution directory, use the CD sequence numbers, release notes information, or information in “Planning the Order of Installation” on page 34 to determine the installation order. If the installation order is wrong, you will see a message to that effect when you launch the installation.

Installing Patch Releases

A patch release contains replacement files that can be installed to modify a particular software product. Installing patches is optional; review the online release notes that accompanied the patch to determine whether to install it or not.

Installing Software for Hardware Upgrades

In some cases, changes to software are required when you install a hardware upgrade on a system that is already in service. As a general rule, adding memory, bitplanes, and disks requires no change to the software, but other hardware upgrades require changes to the installed software. The documentation that accompanies the upgrade explains whether a software installation is necessary for the upgrade.

Installing Accompanying Product Releases

Some hardware upgrades are shipped with a software distribution in the form of a software product release. Use the directions in “Installing Optional Software Products” to install this type of software distribution. If the software distribution contains installation tools, you must use them.

Reinstalling the Same Software

Some hardware upgrades are not shipped with a software distribution, but they require that you reinstall some software after the hardware upgrade is completed. For example, upgrades to a CPU board or the graphics subsystem require a software reinstallation. In such cases, the reinstallation is necessary because the hardware-specific files that were installed for the original hardware are not appropriate for the new hardware.

When reinstalling software after a hardware upgrade, be sure to follow the directions in “Planning the Order of Installation” to determine installation order. Reinstallations require a miniroot installation.

When release notes accompany a product, the product contains an additional *.man* image that contains the release notes. Some reference page subsystems contain reference pages for more than one software subsystem. When you receive your workstation and install a software option for the first time, be sure to check the status of reference page subsystems to verify that the reference pages that you want are installed.

Removing All Software

To remove all installed software from the target system, you must perform a miniroot installation; you cannot remove all software during a live install, since the miniroot is needed to supply the functions that sustain the system until critical software is replaced. For specific instructions on removing software, see *IRIX Admin: Software Installation and Licensing*.

Troubleshooting Software Installations

If you encounter problems while installing software, use the following checklist to identify the cause. More information about troubleshooting software installations can be found in *IRIX Admin: Software Installation and Licensing*.

- If you received an error message, what type of error was it?

There are generally three types of errors in installations: *fatal* (caused by hardware failures or lack of operating system resources), *error* (caused by a command that was unable to execute), and *warning* (caused by a condition that might cause problems later). Each error message should be accompanied by additional descriptive information.

- Did the error message refer to an installation conflict?

Installation conflicts occur when there are unsatisfied product dependencies or when incompatible product are selected for installation. Conflicts can arise because an installed product depends on another product that is not installed, a required subsystem is not marked for installation, or two products are incompatible.

In resolving conflicts, you will often need to make changes in the software you choose to install. In some cases you may need to obtain additional required software.

Software Licenses for Your Work Group

Overview of Software Licensing and Administration

Software licensing and administration are primarily concerns of the administrator. Users do not, for the most part, need to be concerned with issues in this area. Work group computing brings a great deal of flexibility to the setup and distribution of application software. In a work group environment, the administrator can choose to tailor the work group software configuration to the particular needs and resources at hand. The administrator has choices in two major areas: where software applications are stored and where they are executed.

In the typical personal computing environment, applications are stored on each individual machine. While this configuration is also possible in work group computing, the administrator can also choose to keep all application software on a server, or to centralize some applications and distribute others. The considerations in making decisions about where to store software include

- the location of available storage space
- the number of users of a particular application
- the expected frequency of maintenance and updating of the applications
- the economics of the applications needed by the work group

In a work group computing environment, administrators can also configure the distribution of processes throughout the work group. In contrast to a personal computing environment, in which applications needed by a user must always execute on the user's machine, applications in work group computing can be set up to execute on a server. The considerations in making decisions about the distribution of processes include

- the availability of processing cycles across the work group
- the degree to which the needs of work group users for CPU-intensive processes change over the short and long term
- the available bandwidth of the work group network

Central to the configuration of work group application software is the licensing of that software. In a work group environment, licensing also offers a number of flexible options. For a given application, licensing can be configured centrally, on a license server, or it can be distributed to individual workstations. This is completely separate from the issues of where the applications are stored and executed. Thus the administrator has many degrees of freedom in configuring the software/licensing architecture of the work group. The considerations in making decisions about licensing configuration include:

- ease of maintenance and updating of software licenses
- the variation in the number of users needing a particular application simultaneously
- the licensing arrangements available from application software vendors

What Is a License?

A software license is a collection of information that includes at least a license password and a vendor ID. This information authorizes the use of a licensed software product for a period of time (days or years) on a particular number of systems (workstations or servers) at any one time. The use of a licensed software product can be restricted to a particular system, open to the entire work group, or anything in between.

When you obtain a software product, it generally comes with a *license information sheet*. The license information sheet comes in a variety of forms; it might accompany the product, or you might receive it separately by mail, FAX, or e-mail. Be sure to store license information sheets in a safe place. License information sheets include information about the vendor and the product, as well as the license password and duration. You need the information on the license information sheet in order to install and use the software product.

How Licensing Works

There are two types of software licenses: *nodelocked* and *concurrent*.

A concurrent license is sometimes referred to as a floating license. When a concurrent license is installed on a network server, licenses can be made available to anyone running the application on a client system sharing the same network. The number of concurrent users is managed by the license server and limited to a pre-determined number of users.

Generally, any user in the work group has access to a concurrent license from any workstation.

A nodelocked license, in contrast, ties software to a particular host system. That is, when an application has a nodelocked license, it can run only on the machine on which it is installed.

A work group can (and often will) use both concurrent and nodelocked licenses. In addition, an application itself can have both a nodelocked and concurrent license.

License Groups for Concurrent Licensing

A *license group* is a set of systems that use a set of Global Location Brokers and Network License Servers. A licensed application running on a system that is part of a license group requests licenses only from Network License Servers that belong to its license group.

Systems on a network can be divided into one or more license groups and systems on different segments of a multi-segment network can belong to the same license group, but each system can belong to only one license group. Systems belonging to a license group get licenses only from Network License Servers in the license group, and the Network License Servers in a license group do not give licenses to systems outside of the license group. Global Location Brokers and Network License Servers in a license group don't respond to license requests from systems that are not in the license group.

Licensing Software

Silicon Graphics products are licensed by either the NetLS™ or FLEXlm® application (FLEXlm, the newer of the two, is recommended). The *FLEXlm End User Manual* and the *Network License System Administration Guide* are the primary sources of information on setting up and maintaining the licensing scheme at your site.

Nodelocked License Startup

When an application with a nodelocked license starts up, it first checks to see if a license file (usually */var/netls/nodelock*) exists and there is a valid license entry for the application. If there is a valid license, the application runs.

If the nodelock file doesn't exist or there is no valid entry for the application, then the application attempts to run using a concurrent license.

Concurrent License Startup

When an application using a concurrent license starts up, it checks the *Network License Server* for an available license. If a license is available, it is granted, and the application runs. When the user exits the application, the application releases the license.

In checking the Network License Server, an application uses a service called a *Global Location Broker*. A Global Location Broker is a system on the network that maintains a database (the GLB database) containing the address(es) of the Network License Servers on the network, and a list of the vendors that have product licenses at each server.

Any network on which concurrent licenses are used must have at least one Network License Server and Global Location Broker. Large networks may have many Global Location Brokers and Network License Servers. Even in a work group environment, you may want to consider having more than one. For more information, see “Planning and Setting Up Your Licensing Scheme” on page 44.

Note: Concurrent licensing requires a network, while nodelocked licensing can be accomplished on a standalone (non-networked) machine.

Planning and Setting Up Your Licensing Scheme

It is the responsibility of the system administrator to create a software licensing scheme for the work group. The overall steps in planning and setting up a software licensing scheme are as follows:

- Decide whether to implement license groups.
- Set up license groups, if any.
- Choose systems to be Global Location Brokers and Network License Servers.
- Set up Global Location Brokers and Network License Servers.

Choosing Whether to Implement License Groups

License groups are typically used to restrict licenses to a particular set of systems on a network. A license group can include systems on different segments of a multi-segment network. Each system can belong to only one license group. You would normally implement license groups only if your work group consists of subgroups of users who need access to specific applications, and definitely do not need access to applications needed by other subgroups. Because work groups tend to be small and users generally need the same or similar applications, work group environments do not typically implement license groups.

Setting Up License Groups

The first step in setting up a license group is determining the systems that will be included in a license group. This includes the Global Location Brokers and Network License Servers for the group. After the systems are identified, setting up the group involves creating a *glb_obj.txt* file in the */var/ncs* directory on each system in the group. For detailed instructions, see the *Network License System Administration Guide*.

It is also possible to set up license groups manually after Global Location Brokers and Network License Servers are set up. Because manual setup is significantly more difficult, it is recommended that you set up any license groups for your work group first, before setting up Global Location Brokers and Network License Servers.

Choosing Systems to be Global Location Brokers and Network License Servers

Use the guidelines below to choose systems to be Global Location Brokers and Network License Servers. Keep in mind that a single system can be both a Global Location Broker and a Network License Server, and can run licensed applications.

- When a concurrent license has been created for a specific system, that system must be a Network License Server. The system is listed on the license information sheet; the target, sysinfo, or serial number items identify it.
- The systems that you choose for Global Location Brokers and Network License Servers should be stable; they should not be systems that are frequently taken down or unavailable. Good choices are systems that are already designated “server” systems within the network.
- Systems that are Global Location Brokers and/or Network License Servers must be able to use TCP/IP to communicate with each other and with systems running licensed applications.

- If your network uses routers, each network segment where licensed applications are run should have at least one Global Location Broker.
- You should strongly consider setting up two or more Network License Servers. For example, instead of obtaining a license that supports 15 concurrent users and installing it on one Network License Server, obtain three licenses for five concurrent users each and install each license on a different Network License Server. That way, if a Network License Server becomes unavailable, there are still two Network License Servers that can grant permission for ten concurrent users.
- Do not use diskless systems as Global Location Brokers or Network License Servers.

Setting Up Global Location Brokers and Network License Servers

The overall steps in setting up a Global Location Broker and Network License Server are as follows:

- Make sure the license for your software software is intended for the system you want to set up (or that the license can be used on any system).
- Make sure that the necessary software subsystems are installed.
- Set up and run the Network License Server and Global Location Broker.

For detailed instructions, see the *Network License System Administration Guide*.

The Licensing Process

After your licensing scheme is established and your Global Location Brokers and Network License Servers are set up, you can proceed with the process of obtaining, setting up, and installing licenses. Receiving and installing the first license for a software product can be slightly different from installing a replacement license, as detailed in the following figures.

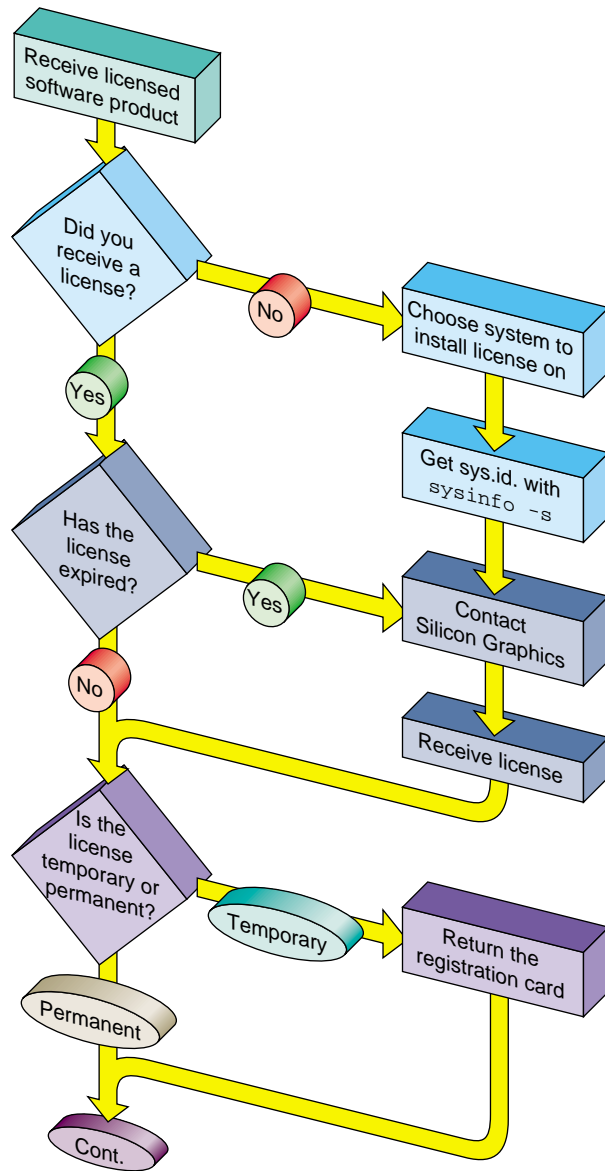


Figure 4-1 Licensing Process for New Software (part 1)

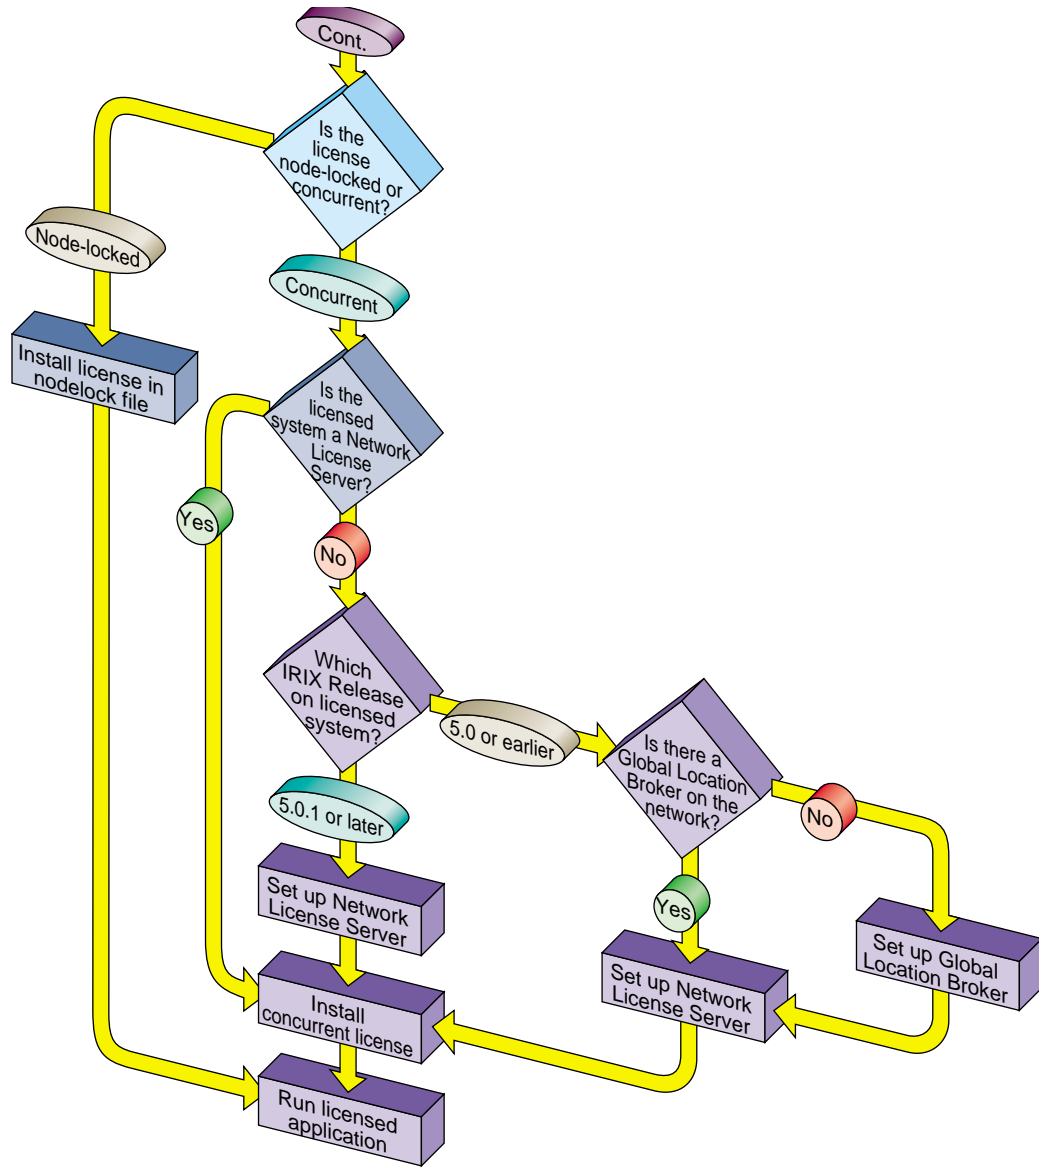


Figure 4-2 Licensing Process for New Software (part 2)

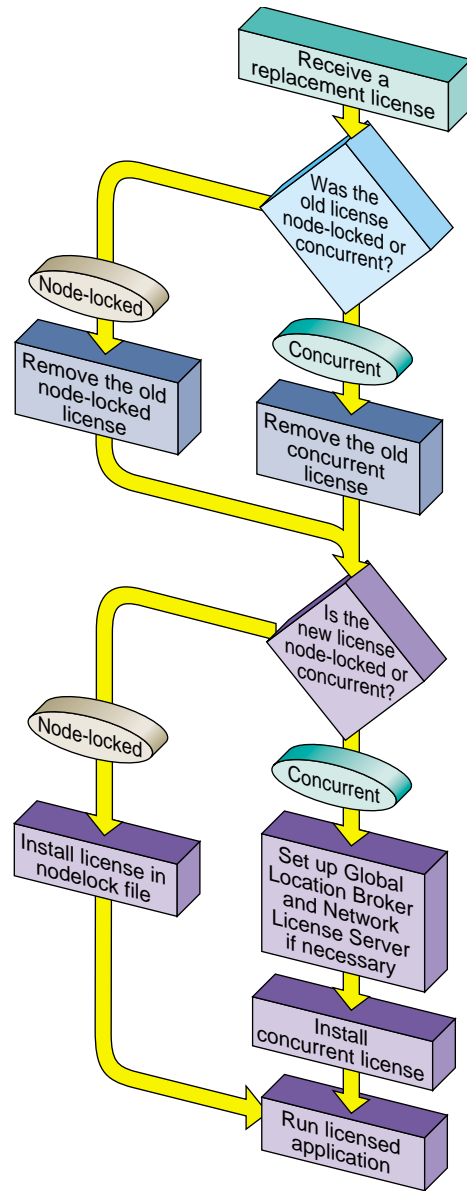


Figure 4-3 Licensing Process for a Replacement License

Maintenance

Software licenses normally require little maintenance and administration. The administrative tasks include the following:

- Verifying that a license was installed properly.
To verify that a software license was properly installed and is functioning, try using an application covered by the license. If the application works, the license is functioning.
- Deleting a license.
To delete a nodelocked license, simply remove the line containing the license and any comments that go with it from the license file (usually */var/netls/nodelock*).
To delete a concurrent license, you run your license administration application as superuser, identify the server, vendor, product, and license you want to delete, and delete it.
For more information about deleting nodelocked and concurrent licenses, see the *Network License System Administration Guide*.
- Monitoring license usage.
You can use your administration software (generally *ls_stat*) to monitor
 - the names of Network License Servers on the network
 - the licensed products at all servers or at selected servers
 - the current users of licensed products and their user ID, host name, group, number of licenses held, and times they started using the products
 - the licenses installed at each Network License Server, the number of active licenses, and their start and expiration dates
 - the use of products: the number of licenses in use, the total number of licenses, and the licenses available

For more information about monitoring license usage, see the *ls_stat* reference page.

- Generating reports on license activity.

You can use your license administration software (generally *ls_rpt*) to generate reports about the activities on a Network License Server. These reports can help you track demand for particular software products. For details about generating license activity reports, see the *ls_rpt* reference page.

- Identifying licenses in a nodelock file.

To locate a particular license in a license file, open the file in a text editor and look for comments that identify the product name associated with each license. Comment lines start with #.

If you cannot find the information you need in comments, try commenting out (by prepending # to each line) all licenses except one. Then run a licensed application. If it runs, its license is the uncommented one. If not, proceed by leaving other licenses uncommented, one by one, until you find the one you want.

To figure out which products the licenses in a license file might apply to, list the installed products by issuing the *versions* command from a shell window. Check the release notes for each product to find out which ones require licenses.

- Installing replacement licenses.

You may need to install a replacement license when the original license expires, when you reconfigure a system or the network in a way that changes a system identification number required by a license, or when you upgrade a software product.

Installing a replacement license generally requires that the original license be deleted. After that, installing the replacement license proceeds exactly as if it were a new license.

Troubleshooting Concurrent Licenses

If you have trouble obtaining a concurrent license for a licensed product, check the following:

- Does a license file contain a nodelocked license for the same product? If so, try commenting out that license.
- Make sure there is at least one Network License Server and Global Location Broker available on the network and that a license is available for the product you want to use.

- Make sure the clock on the user's system agrees (within a minute or so) with the Global Location Broker clock(s), and that the clocks of all the Global Location Brokers also agree.
- Try stopping and restarting all Global Location Brokers and Network License Servers on the network.

For more information about troubleshooting concurrent licenses, see the *Network License System Administration Guide*.

Identifying Your Work Group

About Work Group Identification

One of the chief benefits of work group computing is the ability to use the facilities and information throughout the work group network, rather than relying only on facilities and information on an individual machine. In order to make the resources of the work group available, the first step is to identify those resources, and provide a way to find them simply and easily. The optional Network Information System (NIS) provides that service.

Setting up and administering NIS is an administrative responsibility. In a work group environment, NIS is relatively transparent to users, as it simply provides services they can use as a matter of course.

Introduction to NIS

NIS is a network lookup service that provides a centralized database of information about the network to systems participating in the service. The NIS database is fully replicated on selected systems and can be queried by participating systems on an as-needed basis. Maintenance of the database is performed on a central system.

The purpose of NIS is to make network administration more efficient by reducing the risk of error and the time required to perform redundant file management tasks. For example, maintaining the */etc/hosts* database on a large network might require creating a script to automatically copy the */etc/hosts* file from a central system to all systems on the network. It also requires setting up the appropriate access permissions on each system to enable this file transfer; this is a redundant and time-consuming process. By contrast, on networks using NIS, maintaining the */etc/hosts* database requires modifying a single file, typically */etc/hosts*, on a single system.

NIS can service networks with up to approximately 1000 systems. Larger networks can be organized into multiple NIS service areas, or domains.

NIS Clients and Servers

An NIS *client* is a process running on a system that requests data from an NIS database. An NIS *server* is a process running on a system that provides data from the NIS database. The terms client and server designate both processes and systems: a system is considered a client when requesting NIS data, and it is considered a server when providing NIS data. A system can function as a client and a server simultaneously.

Sometimes client requests are handled by NIS servers running on the same system, and sometimes they are serviced by NIS servers running on a different system. If one NIS server system fails, client processes obtain NIS services from another. In this way, the NIS service remains available even when an NIS server system goes down.

NIS Servers

Every NIS server contains a copy of the NIS database. There are two types of servers:

- A *master server* is a server on which the NIS databases are created and maintained
- A *slave server* contains a duplicate copy of the database.

In most work group environments, a single NIS server is sufficient. However, if slave servers are needed, automatic propagation ensures the consistency of database information between the master server and its slave servers.

NIS Maps

The NIS database is made up of a group of files known as *maps*. Maps are created with NIS tools that convert input files (usually standard ASCII files) to files in database record (dbm) format. Since data in dbm format is faster to find than ASCII data, using dbm files increases NIS performance.

Maps are composed of *keys* and *values*. A key is a particular field in the map that the client must specify whenever it queries the map. Values are attributes of the key returned from the query. For example, a key might be the name of an individual system, and a value might be its Internet address.

NIS Domains

An NIS domain is a collection of systems using the same NIS database. To participate in the NIS service, a system must belong to an NIS domain.

Figure 4-1 shows the basic layout for the systems in Building 1 and a domain called *eng*.

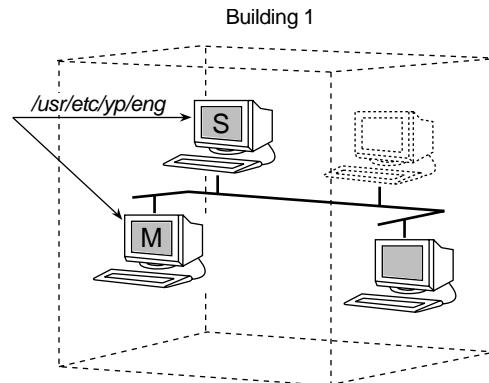


Figure 5-1 Basic NIS Domain

The domain *eng* consists of the master server, one slave server, and three clients. One system on the network does not participate in NIS at this time but may be included in the domain at a later date.

NIS and Internet Domains

NIS domains can be organized to coincide with Internet domain names. In the Internet naming hierarchy, a domain name is comprised of a stream of names separated by dots. For example, if *eng.widgets.com* is an Internet domain, it could also be used as the name of the NIS domain.

Some administrators name their NIS domains with simple names and the Internet domain names separated with dots. For example, the NIS domain name might be *eng*, and the Internet domain name might be *eng.widgets.com*. Using this naming scheme, the two domains can be easily distinguished. Other administrators prefer to have the NIS and Internet domain names coincide. This is strictly a matter of choice; there is no explicit relationship between NIS and Internet domains. However, to avoid confusion, choosing one of these two naming schemes is recommended. (See the *IRIX Admin: Networking and Mail* for details on Internet domains.)

Setting Up NIS

Setting up NIS involves setting up the NIS master server and setting up the clients. An NIS server does not need to be a dedicated system; any system on the network can function as an NIS server.

Setting Up the NIS Master Server

After choosing the system you want to use as the NIS master server for your work group, you must follow four overall steps to set up the NIS master server. For more information about setting up the NIS master server, see the *NIS Administration Guide*.

1. Set the master server's domain name.

Choose a domain name based on your site's configuration. It is possible, but not necessary, for the domain name to coincide with the system's Internet domain name, if any.

2. Build the master maps.

Building the master maps is mostly automatic. You start the process by issuing the appropriate commands to the NIS master server. The server software then proceeds to build maps of names, addresses, and other information.

3. Start NIS on the master server.

You can start NIS on the master server by rebooting the system, restarting the network, or manually starting the NIS software. As soon as NIS is started on the master server, it is available to clients.

4. Test the NIS master server.

The simplest way to check to make sure the NIS master server is functioning properly is to use its services from the same system (which is also an NIS client). One way to check is the command *ypwhich* from a shell window. If the NIS master server is functioning, it should return its own name:

```
# ypwhich  
localhost
```

Setting Up NIS Clients

There are four parts to the procedure for setting up the NIS client. You must repeat these steps for each NIS client you need to set up. You may choose to distribute instructions to work group members and request that they follow the NIS client setup procedure. For more information about NIS client setup, see the *NIS Administration Guide*.

1. Set the domain.

Setting the domain on an NIS client is the same as setting the domain for the NIS master server.

2. Configure NIS on the client.

In most cases, NIS should be set up to start automatically when the client system starts up.

3. Start NIS on the client.

Starting NIS on a client system is the same as starting NIS on the master server. To start NIS you can reboot the client, restart the network, or start the NIS client software manually. Typically, this is done by the following command:

```
# /usr/etc/ybind
```

4. Test the NIS client.

To test an NIS client, you typically issue the *ypwhich* command from a shell window. If the client is functioning correctly, the name of the master server for that client should be returned.

Administering NIS

Administering NIS on your work group network primarily consists of keeping the information in the NIS databases current. This involves these primary tasks, which are covered in more detail in the *NIS Administration Guide*:

- Add new users who join the work group to the NIS database.
- Make sure new users' systems are properly set up to use NIS.
- Maintain NIS passwords for user accounts.
- Create and propagate nonstandard maps, if desired.

NIS maps contain standard types of information. It is possible to store different types of information in NIS. You do this by creating an ASCII file (for example with a text editor) containing the information you want, then creating an NIS map for the file.

To propagate a map you created, you set up a file for the new map in the domain directory of each NIS server on your network. If you use only an NIS master server with no slave servers, you need only store your new map in the domain directory of your NIS master server.

Troubleshooting

If you experience problems with NIS in your work group, check the following:

- Are all NIS clients and servers properly connected to the network?
- Do the clients and servers have the same domain name?

The domain names must match exactly. For example, the domain *widget.com* is not the same as the domain *WIDGET.com*.

- Are your NIS servers overloaded?

When an NIS server is overloaded, the client automatically tries to switch to a more lightly-loaded server. If one is not available, you may need to add additional NIS servers to your network.

- Are your NIS servers running properly?

This will be more of an issue when, in a small work group, you have a single NIS server. In an environment with multiple NIS servers, clients automatically switch to different servers.

If you cannot isolate a problem after making these checks, or if you need additional information about these items, see the *NIS Administration Guide*.

Peripheral Devices

About Peripheral Devices

In work group computing, peripheral devices connected to one workstation or server can be accessible to everyone. The most common peripheral device to be shared is a printer, but other devices, such as optical scanners, and CD-ROM and tape drives, can also be shared. The advantages of making peripheral devices available throughout the work group include the following:

- Maintenance and administration of peripherals such as printers and scanners can be centralized.
- Resources can be more carefully allocated across the work group.
- Services, such as regular backup, can be centralized.

Setting Up Peripheral Devices

It is often the responsibility of the administrator to set up peripheral devices. The first thing to consider in setting up a device is its location. Some peripherals, such as printers, scanners, and modems, tend to be kept in a central location (often nearest to the administrator) for ease of administration, maintenance, and security. Other devices, such as tape drives and floppy drives, may be located at users' workstations.

In determining the best location for peripheral devices, consider these factors:

- How many work group members will typically use the device?
- How often will they need the device?
- What special maintenance does the device require (for example, replacing toner)?
- Does the device require any special security precautions?

The information that follows deals with the overall tasks and considerations in setting up, using, and maintaining peripheral devices for a work group. For detailed instructions about setting up a particular peripheral device, see the documentation that came with it as well as *IRIX Admin: Peripheral Devices*.

Printers

Printers can be connected to individual workstations or to network print server systems. A *print server*, which is a printer connected to network print server system, can be shared by multiple users (each of which is a *print client*).

Work groups often use a print server available to all work group members. It is generally useful to locate a print server centrally, in a location accessible to everyone and convenient for the administrator, who may be responsible for general maintenance and administrative tasks such as replacing toner cartridges and maintaining the print server.

Printers may have parallel, serial, or SCSI interfaces. Any printer, regardless of interface or other hardware characteristics, can be a network print server. Likewise, any system on the network can be a network print server system; it does not need to be a dedicated system. In planning your printer deployment strategy, the volume of print jobs you expect in your work group should play a role in choosing the system or systems you want to use as network print server systems, and in whether those systems are dedicated to that use.

The overall steps in deploying a printer are as follows:

1. Choose a location for the printer.
2. Set up the printer itself.
3. Configure the print server on the print server system.
4. Configure the print server on the print clients.

Configuring a Network Print Server

Print servers must be configured on the print server system before remote clients can configure them successfully across the network.

To configure a print server for use across a network, you simply make sure that the print server system is able to communicate on the network. Then, as superuser, grant permission to each client that will use the print server. Typically, you grant permission to a print client with this command:

```
# addclient client_name
```

The command for adding permission for all remote workstations to use the print server is:

```
# addclient -a
```

For more information about configuring a network print server, see *IRIX Admin: Peripheral Devices*.

Configuring a Print Client

Configuring a print client so a network print server can be used simply involves adding the network print server to the client's Print Manager. You must know the hostname of the print server in order to do this. It is also helpful to know the printer type (for example, a PostScript® printer, line printer, plotter, and so on).

Any printer that has been added to the Print Manager can be designated as the default printer for the client system. In most cases, you should designate a default printer for each client system in your work group.

For more information about setting up printer connections on a print client, see the *Personal System Administration Guide*.

Printer Administration

Administering printing services for a work group generally involves the following tasks:

- Add and remove printers as necessary.
You can use the Print Manager to add and remove printers from client systems.
- Change default printers.

When users' printing needs change, or additional printing resources become available on your network, changing the default printer on some or all workstations may be appropriate. You can use the Print Manager to set as default any printer that has been added to client systems.

- Monitor print server status.

When a print job is sent from a print client to a print server, it can appear to the client that the job is being printed, regardless of its actual status on the print server. To monitor status of a print request on a print server, use the Print Manager, which can detect the true status of a remote print job over the network.

- Cancel print server requests as necessary.

When it is necessary to cancel a print request, you can use the Print Manager. Note that if the print server is not handling a large number of requests, and the document to be printed is short, the job may be printed before you can cancel it.

CD-ROM, Floptical, and Floppy Disk Drives

One of the key concepts in work group computing is that devices connected to one system can be accessible to all systems on the network. The administrator's initial role in regard to CD-ROM, floptical, and floppy disk drives is to determine the appropriate allocation of drives. Depending on the needs of your work group, each system may need its own drive(s), or a single, centrally located drive may suffice.

The overall steps for setting up a CD-ROM, floptical, or floppy disk drive are as follows:

1. Install the drive.

For specific instructions, see the documentation that came with the drive.

2. Mount the filesystem on the drive for use on the local system.

You can use the DiskManager to mount filesystems automatically.

3. If the filesystem will be used on remote systems, export the filesystem.

You must have NFS installed to export a filesystem. See the *ONC3/NFS Administration Guide* for more information on exporting filesystems.

CD-ROM Filesystems

CD-ROM filesystems are always read-only. When you are finished using the filesystem, issue the *eject* command. The filesystem is unmounted and the CD-ROM is ejected from the drive. Any user can unmount and eject a CD-ROM with the *eject* command.

Floptical and Floppy Disk Filesystems

Many different kinds of floptical and floppy disk drives are supported, including the following:

- 720 Kb, 3.5" floppy
- 1.44 Mb, 3.5" floppy
- 20.1 Mb, floptical
- 360 Kb, 5.25" floppy
- 720 Kb, 5.25" floppy
- 1.2 Mb, 5.25" floppy

Silicon Graphics systems automatically determine the format of a floppy disk inserted in your drive, and, if it is a DOS or Macintosh® floppy disk, automatically mounts the filesystem on your default mount directory. Once the filesystem is mounted, you can use typical IRIX commands with it.

You can use a floptical or floppy disk drive like a tape drive for IRIX file transfer. You can use the standard tape archive commands to write files to the floppy disk drive if it is in DOS format.

When you place files on a floppy disk, it is a good idea to make a note on the disk label of the format or the exact command used to place the files on the floppy disk. This makes it much easier for you (and others) to retrieve the files from the floppy disk. Also, whenever possible, change directories to the one that contains the file and place the file on the floppy disk using a relative pathname, rather than specifying the absolute pathname.

Also, be aware that using a floppy disk to transfer files to systems made by other manufacturers may mean that the same tools are not available on the receiving system. The *tar*, *cpio*, and *dd* tools are usually available on all UNIX systems.

Tape Drives

This section covers what you need to know about the tape drives on your workstation or server. The cartridge tape device is used primarily for filesystem backups and data transfer.

For information on backing up data onto tapes, see *IRIX Admin: Backup, Security, and Accounting*. If you are installing a tape drive, see the installation instructions furnished with the hardware.

Almost all workstations are configured with some sort of tape device for making backup copies of your files. Whether you maintain one system or a network of hundreds of workstations, you will eventually have to use and maintain some form of tape drive.

Adding a Tape Drive

To install a tape drive on an IRIX system, follow the hardware installation instructions furnished with your tape drive. Make sure you carefully follow any instructions regarding drive terminators.

If you are adding a tape drive to a system that does not have one, the software configuration is taken care of automatically when the system boots.

Troubleshooting

This section discusses some troubleshooting tips for problems arising with peripheral devices.

Troubleshooting Your Printing System

If you send a print request to a printer and you do not receive any output, use the checklist below to find the problem. You can find additional troubleshooting information in the printer manufacturer's hardware manual, and in *IRIX Admin: Peripheral Devices*

- Is the printer turned on?

Printers do not always indicate clearly if they are turned on. Make sure the printer is plugged into the power socket and the power switch is on.

- Does the printer have paper?

Frequently, printers run out of paper when printing in high volume.

- Is there a paper jam?
Make sure the entire paper pathway is clear of sheets or fragments of paper. Refer to your printer hardware documentation before attempting to put any unusual paper or other media through your printer.
- Is the printer set to the correct baud?
Be sure the baud rate of the printer matches that of the serial port.
- Is the serial cable attached correctly?
Often, reseating the serial cable where it connects to the printer restores correct operation.
- Is the correct cable being used?
The use of the pins in serial cables varies somewhat in different applications. Cables designed for specific hardware may or may not function correctly with different hardware. Check your system Owner's Guide and the documentation supplied with your printer and cable to determine if the cable is correct for your hardware.
- Did you specify the right printer?
If your system has more than one printer, and you wish to send a job to a printer other than the default, remember to explicitly choose the printer you want to use.

Troubleshooting Tape Drives

If you have difficulty accessing a tape drive, use the checklist below to find the problem. You can find additional troubleshooting information in the drive manufacturer's hardware manual, and in *IRIX Admin: Peripheral Devices*.

- If the tape drive is an external unit, does it have power?
Simply powering it on does not cause it to be seen by the computer. The system must be shut down, power cycled, then rebooted.
- During the boot phase, do you see the access light on the tape drive light up at all?
If it doesn't flash at all, chances are the operating system is still not seeing the drive.
- Is the SCSI cabling and termination correct?
If visual inspection shows nothing obvious, try resetting the connectors. Any movement of hardware or cabling must be done with the system powered off.

- Was the tape device's SCSI address changed when other SCSI devices were added to the system?
- Is the tape drive's read/write head clean?

Follow the manufacturer's guidelines for maintenance and cleaning.

This covers the basic problems that administrators experience regarding missing tape drives. See the following reference pages for more information on the commands used in this section: *mt*, *ls*, *hinvo*. For more technical information about tapes, see *mtio*, *tps*, or *xmt* reference pages.

Troubleshooting Tape Read Errors

Often there is a quick and simple fix for an error message that is caused by a tape drive malfunction or the tape itself. Both recoverable and unrecoverable errors can be caused by something as basic as a dirty read/write head, a poorly tensioned tape, or a dropout, which is a physically bad spot on the tape. An EOT message can also mean that there is no data on the tape.

The commands below could help you with some basic tape maintenance and performance functions. Using these commands could either prevent future errors from occurring or help you recover from an existing error condition:

- The *hinvo* command determines which tape drive type is connected to your system.
- The *mt status* command verifies the status of the tape drive and the media.
- The *mt retension* command ensures that consistent tension is applied to the tape within the cartridge.
- The *mt reset* command resets the tape drive and controller. Use this command only as a last resort; in some instances it may result in loss of data. In the case of SCSI tape drives, *mt reset* resets all SCSI devices, including disk drives.

Overview of System Security

System security under IRIX is primarily dependent on system login accounts and passwords. Proper administration, user education, and use of the facilities provided yield adequate security for most sites. Most security breaches are the result of human error and improper use of the provided security features. No extra measures yield more security if the basic features are not used or are compromised by user actions. Also, periodically log in with anonymous FTP to sgigate.sgi.com and look in the directory `~ftp/security` for any security patches for your system.

This chapter deals with maintaining security in a work group environment. Security is up to both administrators and users — it starts with your own local system. Once you have initially established the security of a system, you can expand your secure area to include the network. But until you have local security, there is no point in trying to establish security over a larger area.

The most important concept in security is that security is a dynamic process, requiring that you understand the issues, keep up to date on them, and continually monitor your system and work group. Each individual in the work group must be aware of security and take steps to maintain it. Your responsibility as an administrator is to establish a security policy and begin its implementation.

A great strength of the IRIX system is the ease with which users can share files and data. However, some of the security guidelines presented in this chapter are at odds with easy system access. It is up to the system administrator to balance the needs and concerns of the user community.

There are three levels at which security must be established: individual systems, the local network, and a site connected to external networks such as the Internet. This chapter deals primarily with taking steps to secure an isolated system, but many of these same steps must also be taken before undertaking the more ambitious job of securing a network.

Components of Security

The components of security include physical security of systems, passwords protecting access to systems and resources, and permissions protecting files and directories.

Physical Security

The physical security of the systems in a work group is usually only partially the administrator's responsibility. The administrator's primary role is to educate other work group members in physical security. For the most part, this consists of limiting access to physical systems, peripherals, and storage media. It can also include providing individual systems and the network with protection against power surges, outages, and fluctuations.

Passwords

Managing passwords is also described in *IRIX Admin: System Configuration and Operation*.

A system is most secure if nobody can access the system without an account and password, and if all the passwords on the system are difficult to guess and obtain. Surprisingly, many users choose passwords that are easy for potential intruders to guess, or write their passwords down on paper and leave them near their workstations and terminals.

Some site administrators use the same password for multiple administrative accounts. This is not a good practice. Do not deliberately use the same password for more than one account.

More secure passwords have all of these features:

- They are long (the first eight characters are recognized).
- They use multiple words that are combined or arranged in an unusual manner.
- They contain words from multiple languages, combined in a unique way.
- They are composed of different kinds of characters, such as digits and punctuation

Easily guessed passwords have these characteristics:

- They are short.
- They consist of single words that are in a dictionary.
- They are the same as the account name, or the account name spelled backward.
- They contain the name of the user's department or project.
- They consist of the user's name or initials.
- They contain the license number of the user's car, a spouse or friend's name, the user's home address, phone number, age, or other obvious information.
- They are obvious—for example, "top secret," "secret," "private," "password," "friend," "key," "god," "me," and so on

PROM Passwords

Your system has a facility that allows you to require a password from users who attempt to gain access to the Command (PROM) Monitor. This gives greater control over who may perform system administration tasks.

Traditionally, if an intruder gains access to your system hardware, there is little you can do to maintain system security. In the simplest case, the intruder switches off the system, then turns it on again, and instructs the system from the console to boot a program other than your usual operating system. Alternatively, the intruder could simply remove the hard disk from your system and install it on another system and read your files. While there is nothing you can do with system software to prevent physical theft of the hardware, you can limit the ability of intruders to boot their programs or to otherwise damage your system at its lowest levels with a PROM password.

You can reset the PROM password on most systems as long as you have the root password. If you cannot successfully reset the PROM password, you must remove the PROM or a jumper from your CPU board. See the *Owner's Guide* for the system for information on this procedure.

Second (Dialup) Passwords

If a system requires additional protection, you can establish a system password. If you do this, users who log in on specific ports (ttys) are prompted for a system password in addition to their account passwords.

System passwords are normally used only on dialup lines and are often referred to as dialup passwords. You can use them on standard lines, but this is usually not necessary.

Creating a Shadow Password File

Password files are generally accessible to any user of a system. Although the file is encrypted, if a copy of the file can be made, someone could attempt to decrypt it. You can prevent this type of attack by using a “shadow” password file. A shadow password file is a copy of the standard password file that is not accessible by non-privileged users.

The shadow password file is called */etc/shadow*. Once shadow passwords have been initialized, the password field in each */etc/passwd* entry is replaced by an “x” character.

Note: Shadow passwords work differently with NIS. See the shadow(4) reference page for details on the use of shadow passwords with NIS.

Password Aging

You can use the password aging mechanism to force users to change their passwords periodically. It also prevents a user from changing a new password before a specified time interval. You can also force a user to change his or her password immediately.

Password aging does not provide as strong protection as it might seem; most users simply choose two passwords instead of one. This is because, when password aging is enforced, most users alternate between two passwords that they find easy to remember rather than inventing new passwords every time their old ones expire.

You can set passwords to expire after any length of time, from a day to several months or longer.

Note: Password aging is not supported for NIS entries.

Checking the Password File

From time to time, you should scan the password file on your system (and optionally, perform this procedure on users’ systems as well). You can check the password file for completeness and consistency of information. If unauthorized access has been attempted, this type of checking may detect it. You can validate

- the number of fields in each entry
- the login name
- the user ID number
- the group ID number
- the login directory
- the executed program

To check a password file, you typically use the *pwck* command.

File and Directory Permissions

Be conservative when establishing or changing permission bit settings on all files and directories. The safest settings do not allow write access, but where this is not possible, it may be possible to limit write access to the owner of the file or directory, or at least just to the owner and the group.

The following files and directories are universally available for read and write access on IRIX as shipped. Depending on your site requirements, you may wish to change the permissions on these files to be more restrictive.

Caution: Restricting permissions on historically open directories, such as */tmp*, */usr/tmp.O*, and */var/tmp* (linked to */usr/tmp*), can cause serious malfunctions in many programs, applications, and system utilities that write temporary files on behalf of users in these directories. Below is a partial list of such directories:

- */tmp*
- */usr/demos/.xsession*
- */usr/Insight/tmp*
- */usr/Insight/tmp/ebtpriv*
- */usr/Insight/tmp/ebtpub*
- */usr/Insight/tmp/install.insight.log*
- */usr/lib/emacs/maclib*
- */usr/lib/showcase/fonts*
- */usr/lib/showcase/images*

- */usr/lib/showcase/models*
- */usr/lib/showcase/templates*
- */usr/tmp.O*
- */var/spool/locks*
- */var/spool/uucppublic*
- */var/tmp*

Security Checklists

For additional information about security, see *IRIX Admin: Backup, Security, and Accounting*. In attending to security of systems, networks, and information, both administrators and users should be aware of the following:

- Anyone with physical access to a computer can simply take it or take its disk drives(s).
- The same caveat applies to backups of the system; anyone with physical access to backup tapes can gain access to any information stored on them.
- Common-use accounts are a potential security hole. An example of a common-use account is one that is shared by all members of a department or work group. Another example is a standard “guest” account on all the workstations at a site. This allows all users at the site access to different workstations without requiring specific accounts on each workstation.

A pitfall of common-use accounts is that you cannot tell exactly who is responsible for the actions of the account on any given workstation. Another risk is that anyone trying to break into workstations at your site will try obvious account names such as guest.

Common-use accounts can be helpful, but be aware that they can pose serious security problems. Needless to say, common-use accounts that do not have passwords are especially risky.

- Accounts that are no longer used should be either locked or backed up and removed, since unused accounts can be compromised as easily as current accounts.

Also, change critical passwords, including dialup passwords, whenever anyone leaves the organization. Former employees should not have access to workstations at the site.

- Systems with dialup ports should have special dialup accounts and passwords. This is very important for sites that have common-use accounts, as discussed above.

Even with this added precaution, do not store sensitive data on workstations that have dial-up access.

- If your site allows access to the Internet network (for example, using *ftp*), take precautions to isolate access to a specific gateway workstation.
- Discourage use of the *su* command unless absolutely necessary. The *su* command allows a user to change his or her user ID to that of another user. It is sometimes legitimately necessary to use *su* to access information owned by another user, but this presents an obvious temptation: the person using *su* to switch user IDs must know another person's password and therefore has full access to his or her account.

Note: The file */var/adm/sulog* contains a log of all successful and unsuccessful attempts to use the *su* command (if it is enabled in */etc/default/su*).

- Be sure that system directories such as */* (root), */bin*, */usr/bin*, and */etc* and the files in them are not writable except by the owner.
- If you must leave your console, workstation, or terminal unattended, log off the system. This is especially important if you are logged in as root. Also, refer to the *xlock* reference page for information on locking your local X display.
- Sensitive data files should be encrypted. The *crypt* command, together with the encryption capabilities of text editors (for example, *ed* and *vi*), provides some protection for sensitive information.

Security for Administrators

Site administrators should be aware of the following:

- Permissions for directories and files should be set to allow only the necessary access for owner, group, and others. This minimizes the damage that one compromised account can cause.
- There are several ways accounts and passwords protect the system:
 - By requiring users to log in with specific accounts, you can determine who is responsible for specific actions on the system.
 - Using the IRIX system of file permissions, users can keep data reasonably secure. Other users on the system are less likely to view confidential material accidentally.

- If all accounts have passwords, the chance of an unauthorized person accessing the system is greatly reduced. However, the possibility of unauthorized access increases if users are lax about changing their passwords regularly and choosing good passwords. The section “Passwords” on page 68 describes how to choose good passwords.
- All active accounts need passwords, which should be changed regularly. Do not use obvious passwords, and do not store them online in “plain-text” format. If you must write them down on paper, store them in a safe place.
- Make sure that each user’s home account, and especially the shell-startup files *.profile*, *.login* and *.cshrc*, are writable only by that user. This ensures that “Trojan horse” programs are not inserted in a user’s login files. (A Trojan horse program is a file that appears to be a normal file, but in fact causes damage when invoked by a legitimate user.)
- Safeguard and regularly check your network hardware. One possible way to break into computer systems is to eavesdrop on network traffic using physical taps on the network cable. Taps can be physical connections (such as a “vampire tap”) or inductive taps.
- Run networking cable through secure areas and make sure it is easy to examine regularly. Create and maintain a hard-copy map of the network to make it easier to spot unauthorized taps. Another way to make this sort of attack less likely is to use fiber-optic (FDDI) network hardware, which is much more difficult to tap.
- Use only that software that is provided by reputable manufacturers. Be wary of programs that are distributed “publicly,” especially already-compiled binaries. Programs that are available on public bulletin board systems (as opposed to BBSs run and sponsored by vendors) and on public computer networks could contain malicious “worm” routines that can violate system security and cause data loss.

Sharing Documentation

Overview and Benefits of Sharing Documentation

The documentation for your Silicon Graphics systems is available online. Because the documentation is quite extensive, it makes sense to store the online files centrally, rather than replicating the files on each system in the workgroup. The administrator generally sets up and maintains the shared documentation for the work group. You can share documentation throughout your work group by using the IRIS InSight™ system on a DynaWeb™ server. DynaWeb enables you to gain access to the shared documentation by using a browser (for example, Netscape Navigator™).

IRIS InSight and DynaWeb

A DynaWeb server makes IRIS InSight books available to World Wide Web clients. DynaWeb dynamically converts the InSight Standard Generalized Markup Language (SGML) tagged text and graphics into HyperText Markup Language (HTML) form.

This section explains the components of the DynaWeb server including the programs used to administer it, the directory structure of the DynaWeb Server, and the features of the DynaWeb server that give the individual administrator great flexibility configuring the server.

How DynaWeb Works

DynaWeb has two major components: a server and an SGML-to-HTML converter.

Like any Web server, a DynaWeb server listens for HTTP requests from clients. Each request contains a URL that identifies the server and a particular block of data such as a home page, a section of text, or an image.

IRIS InSight books are binary files that contain SGML-tagged text and graphics. The SGML-to-HTML converter transforms each request for all or part of an InSight book into an HTML-tagged text stream. If a large block of text is requested, the converter dynamically builds a table of contents (TOC) for the text.

These components work together to handle requests from Web clients.

The server can simultaneously handle up to 256 incoming requests from Web clients. Whenever a Web client requests data from a DynaWeb server, the browser sends an HTTP request to the server specifying the information to be sent. On receiving the request, the server establishes a connection with the client and passes the data request to the DynaWeb SGML-to-HTML converter. The information is transmitted back to the client.

The connections are short lived, lasting only long enough for the server to process the requested data and send it to the client. After sending the requested information, the server terminates its connection.

DynaWeb Directory Structure

The DynaWeb directory structure helps to isolate other parts of the filesystem from Web clients. The directory */usr/lib/Insight/dweb/serverroot* is the highest point in the DynaWeb directory tree that any Web client can access and is called the DynaWeb root directory. The paths of all files accessible to Web clients are defined relative to the DynaWeb root directory. In this way, the root directory helps control the part of the host filesystem that is accessible to Web clients.

Collections of books are added to the DynaWeb directory structure by linking or copying directories to the */usr/lib/Insight/dweb/serverroot* directory.

Collections

A DynaWeb collection is an IRIS InSight bookshelf containing one or more InSight books that a DynaWeb server offers to Web clients. Each collection is a directory in the DynaWeb *root* directory and each book in a collection is a subdirectory of the collection directory. Collections are added to the DynaWeb root directory in three ways:

- Each time the DynaWeb server is started, it creates links (if they don't already exist) from the *servroot* directory to directories that contain InSight books.
- You can create symbolic links from the DynaWeb *root* directory to a collection directory.
- You can copy a collection to the DynaWeb *root* directory.

InSight book collections are automatically made available to the DynaWeb server by means of symbolic links in the DynaWeb root directory. Using symbolic links in this way simplifies book maintenance and DynaWeb upgrades.

Because the DynaWeb server collects information about available collections and books only when it starts up, changes to the collections and books that happen while the server is running are not displayed to client users. When collections are added, books are added, or books are removed, the DynaWeb server must be restarted so that it displays the available collections and books.

You can limit access to the documentation server by setting up DynaWeb to use HTTP authentication. In that case, users must enter a password to view the available collections and books. For more information, see the *IRIS InSight DynaWeb Administrator's Guide*.

Home Page

A home page is an HTML document that welcomes client users to the DynaWeb *root* directory on the server, introduces your Web site to them, and provides at least one predefined browsing choice for accessing information offered by the server. DynaWeb supplies a default home page, which can be replaced by another home page.

Publisher's Page

A publisher's page is an optional HTML document that publishers can use for any purpose, such as presenting forms to collect data from client users, or providing a list of other sites of interest. For example, a publisher's page might enable readers to submit evaluations of documents or notes about improvements or upgrades.

Access and Error Logs

DynaWeb supports an access log file and an error log file. By default, the server logs information to both files. However, you can turn off either or both kinds of logging.

You can extract statistics such as the total number of accesses, the number of requests, and the type of information requested.

CGI Scripts

DynaWeb supports Common Gateway Interface (CGI) scripts. CGI defines an interface that lets your server run scripts. Some sample scripts are supplied with DynaWeb. For additional information, including a complete description of the Common Gateway Interface, refer to <http://hoohoo.ncsa.uiuc.edu/cgi/>.

Clickable Graphics

DynaWeb supports hot clickable graphics that link to different destinations when you click different parts of the graphic.

Setting Up DynaWeb for Your Work Group

This section describes the considerations and overall process of setting up DynaWeb for your work group. For more detailed instructions, see the *IRIS InSight DynaWeb Administrator's Guide*.

DynaWeb Setup Options

DynaWeb can run either as a CGI script or as a process. If you are concerned about security, you may wish to run DynaWeb as a CGI script piggybacked on another server, utilizing the SGML-to-HTML and organizational capabilities of DynaWeb while retaining the extra security provided by the front-end server.

If you set up DynaWeb to run as a process, it can be set up so that it appears to users as multiple servers. Each "server" can have a different configuration, for example, different collections of books or server pages in different languages. Each configuration is described by a configuration file and is associated with a port number. Users select the configuration they want using hypertext links or by specifying a port number in a URL.

The choices you make in setting up DynaWeb can affect the methods users employ to access books and collections. The following methods can be used regardless of whether you set up DynaWeb to run as a CGI script or as a process:

- Create a hypertext link from your home page to the DynaWeb home page.
- Specify the URL of the DynaWeb home page or the DynaWeb collections page to your Web browser.
- Install the Silicon Graphics What's New product.

If you set up DynaWeb to run as a CGI script, this method can be used to access books and collections:

- Specify the URL of the parent server to your Web browser.

DynaWeb Setup Process

Following the initial setup of DynaWeb to run as a CGI script or as a process, follow these steps to make IRIS InSight book collections accessible to your work group:

- Set up collections in standard locations
IRIS InSight book collections in the standard locations, */usr/share/Insight/library/*/**, are automatically made available to the DynaWeb server when it starts up. InSight book collections in nonstandard locations and collections of DynaText books can be made available as well.
- Optionally, set up collections in non-standard locations.
You can set up collections in non-standard locations and collections of DynaText books in two ways: by using symbolic links or by setting up a collection as a directory. Using symbolic links makes book maintenance and DynaWeb upgrades simpler, but introduce potential security risks, particularly if the linked directory contains additional links.
- Optionally, enable HTTP authentication.
When the DynaWeb server is installed, authentication is inactive. To enable user name and password checking, uncomment the line that starts `DWEB_PASSWD_FILE =` in the configuration file */usr/lib/Insight/dweb/data/config/dynaweb.cfg.tmpl*, and add the absolute pathname of the password information file as the value of `DWEB_PASSWD_FILE`.

- Optionally, make other changes to the configuration parameters.

Ordinarily you do not need to change the parameters manually in the configuration file. Several values (such as hostname and server port) are defined during installation. Most other values have defaults. The configuration parameters and the format of the configuration file are detailed in the *IRIS InSight DynaWeb Administrator's Guide*.

Note: Collections you add to DynaWeb and changes you make to the DynaWeb configuration parameters do not take effect until the server is restarted.

Troubleshooting DynaWeb

This section lists possible causes and solutions for some problems that you may encounter when administering a DynaWeb server.

Directories

- You cannot find a DynaWeb directory using the pathname specified in a DynaWeb guide.
A directory pathname without a leading slash indicates that the path is relative to the DynaWeb home (root) directory.
- Clients are unexpectedly accessing directories outside of the DynaWeb root directory tree.

Make sure that symbolic links exist only where expected. Web clients can navigate downward within any linked directory. Moreover, if a linked directory is, itself, linked to another directory, Web clients can jump to and move downward from there.

Home Page

- The server does not display your site-specific home page.

The filename of the home page is set by the `DWEB_HOMEPAGE` configuration parameter. To use a different filename for the home page, you must assign your own filename as the value of `DWEB_HOMEPAGE`.

- The hypertext link to your collections fails to bring up your collection list.

Ensure that the `XREF` value in the link matches the value of the `DWEB_COLLECTMAGIC` configuration parameter.

- The server cannot find your home page after you change its name in the server configuration file.

Check the file and its pathname.

The values of resource parameters must contain a leading slash (for example, `/docs/my-homepage.html`). This slash is necessary, because the server appends the value of file parameters to the value of `DWEB_ROOTDIR` to produce the full directory path.

Client Problems

- A client can't access the DynaWeb server.
 - Make sure the server is responding.
 - Check the configuration file for errors.
 - Make sure the file `/usr/lib/Insight/dweb/data/config/dynaweb.cfg` exists.

- A button does not appear on the button bar.

Check the server configuration file to ensure that the button assignment is not disabled (by the insertion of a `#` at the start of the line). Make sure that the icon file is present in the `servroot/icons` directory and uses the same name as defined in the server configuration file.

Note: An unavailable icon causes no problem for the server.

- A client displays a browser-specific icon, instead of the correct server icon.

Make sure the `.gif` file for the icon exists and is not corrupted.

- The server cannot find an icon after you change its name in the configuration file.
Make sure the values of resource parameters contain a leading slash (for example, */icons/button_up.gif*).
- Users cannot find the button bar.
If the button bar is at the bottom, and the returned information is longer than the document view, users must scroll downward before the button bar appears. In the middle of a block of text, the button bar is out of view (at either the top or bottom of the text).
- A bookshelf is missing or a removed bookshelf or book still appears.
Make sure you restarted the DynaWeb server after adding or removing a book or collection.
- A user complains that the server is acting weirdly.
Verify whether the user is specifying either of the following parameters in URLs: DWEB_NAVHINTS or DWEB_REDIRECT. The server uses these parameters to pass state information to itself. Setting them in a URL causes the server to behave in unpredictable ways.

Authentication

- After you supply your ID and password and they have been validated by the server, when you switch books, collections, or other documents, you are again prompted for your information.
Make sure all URLs match the DWEB_HOST parameter. Otherwise the server thinks you are not authenticated and tries to validate you again. This will happen every time you link to a URL that contains a different hostname. For example, *mycomputer:88* does not match *mycomputer.mycompany.com:88*.

Tell Us About This Manual

As a user of Silicon Graphics products, you can help us to better understand your needs and to improve the quality of our documentation.

Any information that you provide will be useful. Here is a list of suggested topics:

- General impression of the document
- Omission of material that you expected to find
- Technical errors
- Relevance of the material to the job you had to do
- Quality of the printing and binding

Please send the title and part number of the document with your comments. The part number for this document is 007-3484-001.

Thank you!

Three Ways to Reach Us

- To send your comments by **electronic mail**, use either of these addresses:
 - On the Internet: techpubs@sgi.com
 - For UUCP mail (through any backbone site): *[your_site]!sgi!techpubs*
- To **fax** your comments (or annotated copies of manual pages), use this fax number: 415-965-0964
- To send your comments by **traditional mail**, use this address:

Technical Publications
Silicon Graphics, Inc.
2011 North Shoreline Boulevard, M/S 535
Mountain View, California 94043-1389